Aalto University School of Science
Department of Computer Science

Seminar on Network Security and Internetworking

Spring 2015

Mario Di Francesco, Sanja Šćepanović (eds.)

Tutors:

Çağatay Ulusoy, Deng Yang, Vu Ba Tien Dung, Jiang Dong, Jukka K. Nurminen, Keijo Heljanko, Mehrdad Bagheri Majdabadi, Mario Di Francesco, Manik Madhikermi, Matti Siekkinen, Nguyen Trung Hieu, Otto Huhta, Sanna Suoranta, Sanja Šćepanović, Sakari Luukkainen, Sandeep Tamrakar Thomas Nyman, Zhonghong Ou

Aalto University School of Science
Department of Computer Science

Aalto-yliopisto
Aalto-universitetet

Distribution:
Aalto University School of Science
Department of Computer Science
P.O. Box 15400
FI-00076 Aalto
Tel. +358-9-470 23228
Fax. +385-9-470 23293

Preface

The *Seminar on Network Security* and *Seminar on Internetworking* are Master's level courses in computer science at Aalto University. These seminar series have been running continuously since 1995. From the beginning, the principle has been that the students take one semester to perform individual research on an advanced technical or scientific topic, write an article on it, and present it on the seminar day at the end of the semester. The articles are printed as a technical report. The topics are provided by researchers, doctoral students, and experienced IT professionals, usually alumni of the university. The tutors take the main responsibility of guiding each student individually through the research and writing process.

The seminar course gives the students an opportunity to learn deeply about one specific topic. Most of the articles are overviews of the latest research or technology. The students can make their own contributions in the form of a synthesis, analysis, experiments, implementation, or even novel research results. The course gives the participants a personal contacts in the research groups at the university. Another goal is that the students will form a habit of looking up the latest literature in any area of technology that they may be working on. Every year, some of the seminar articles lead to Master's thesis projects or joint research publications with the tutors.

Starting from the Fall 2014 semester, we have merged the two alternating courses, one on security and one on internetworking, into one seminar that runs on both semesters. Therefore, the theme of the seminar is broader than before. All the articles address timely issues in security and privacy and networking technologies. Many of the topics are related to mobile and cloud computing and to the new applications enabled by ubiquitous computing platforms and network connectivity.

These seminar courses have been a key part of the Master's studies in several computer-science major subjects at Aalto, and a formative experience for many students. We will try to do our best for this to continue. Above all, we hope that you enjoy this semester's seminar and find the proceedings interesting.

Mario Di Francesco
Professor

Sanja Šćepanović
Editor

Table of Contents

Design trade-offs for building a real-time Big Data system based on Lambda architecture

Hussnain Ahmed
Student number: 281557
hussnain.ahmed@aalto.fi

Abstract

Major Big Data technologies, such as MapReduce and Hadoop rely on the batch processing of large data sets in the distributed parallel fashion.The latencies due to batch processing techniques are unsuitable for use in real-time or interactive applications. Real-time stream processing engines can process data in real-time but lack the capacity for handling large volumes of data. Lambda architecture has emerged as a powerful solution to provide the real-time processing capability over large volumes of data. Lambda architecture combines both batch and stream processing, working together in a single system transparent to the end user. It provides the basic guidelines of a construct for such data system but allows flexibility in using different components to achieve real-time Big Data processing capability. In our study, we provide a working Lambda architecture implementation while discussing the underlying trade-offs for the design of real-time data systems based on this architectural paradigm.

KEYWORDS: Big Data, analytics, lambda architecture, streaming, distributed computing

1 Introduction

Ubiquitous computing, availability of the fast and mobile Internet and the phenomenal growth in the use of social media have generated a major surge in the growth of data. Advancements in distributed parallel computing have become the major source for balancing this punctuated equilibrium. A strong collaboration between industry and open source software communities has resulted in new programming models and software frameworks, such as MapReduce and Hadoop, to handle Big Data in distributed parallel fashion. A generation of new tools and frameworks is emerging within the same ecosystem as building blocks to enable end-to-end Big Data platforms. The Hadoop framework provides scalability, reliability and flexibility to handle large volumes of data in a variety of different formats. However, Hadoop is designed for distributed parallel batch processing. It can run batch computations on very large amounts of data, but the batch computations have high latencies [14]. Many real life business applications of Big Data, such as web analytics, online advertisements, Internet of things, social media analytics, and operational monitoring, require real-time processing of large streams of data. The problem is that the data processing latencies can lower the efficacy of such applications.

Byron Ellis, 2014 differentiates streaming data from the other types of data on the basis of three major characteristics, i.e. the "always on always flowing" nature of the data, the loose and changing data structures, and the challenges presented by high cardinality dimensions [6]. These three characteristics also dictate the design and implementation choices to handle the streaming data. Another important requirement for such systems is their ability to analyze the live streaming data along with the large volumes of stored historical data. The final outputs of such data systems are usually the combined results, derived from streaming and stored data processing. Recently we have seen some new tools and techniques to manage such data processing. We have already mentioned Hadoop and its ability to batch process Big Data in the distributed parallel manner. Hadoop 2.0 (YARN) and in-memory distributed batch processing within Apache Spark framework was introduced to reduce the data processing latencies. Similarly, tools such as Apache Storm have become very popular as the answer for distributed processing of data streams. Various other tools for functional components such as data collection, aggregation, and distributed database systems are also available. However, significant efforts are required to make appropriate architectural choices to combine these components in the form of a real-time Big Data analytics platform.

Lambda architecture [13] is a design approach that recommends combining the distributed batch processing with stream processing to enable real-time data processing. This approach dissects data processing systems into three layers, i.e. a batch layer, a serving layer and a speed layer [1]. The stream of data is dispatched to both the batch and speed layers. The former layer manages the historical data sets and pre-computes the batch views. The serving layer indexes the batch views in order to serve queries at a low level of latencies. The Lambda architecture can be implemented using various combinations of the available tools, such as Hadoop File System (HDFS), Apache Hive and Apache Spark for batch view generation, Apache Kafka and Apache Storm in the speed layer and HBase in the serving layer. Although Lambda architecture provides a working model to facilitate real-time Big Data analytics, there are some weaknesses associated with it. As highlighted by Jay Kreps, 2014 [10] the real cost of this approach is to maintain the same business logic in two different systems: once in the batch views and then in the speed layer. Such approach adds operational overheads and raises questions regarding the efficiency of this approach.

Our study focuses on Lambda architecture in detail. The key element is the continuous delivery of real-time Big Data analytics using distributed parallel computing. We discuss the design choices for building end to end real-time Big Data analytics pipelines, for understanding the underlying trade-offs of distributed computing, and presenting a working model of lambda architecture.

2 Lambda architecture

Lambda architecture provides an architectural paradigm for real-time Big Data systems. Nathan Marz presented Lambda architecture in his book "Big Data: Principles and best practices of scalable real-time data systems" [14]. The main idea behind Lambda architecture is to run ad hoc queries on large data sets [2]. The large data sets may contain large amounts of the stored data as well as the real-time streaming data. Lambda architecture simplifies this by precomputing the data views. The queries run on these precomputed views as an abstraction of the full data set [14].

$$query = function(all\ data)$$

The overall data system consists of three layers i.e. batch, speed and serving layer. The batch layer runs batch processing jobs periodically on the large historical data sets in order to create the batch view(s). From a data ingestion point of view, the streaming data appends to the historical data as a master copy of the data. For instance, if the batch computations are computed on a daily basis, today's data will be stored with past data and the next batch that will run tomorrow will also include today's data to generate the most recent batch view. Nathan Marz describes this in the equations as below [14].

$$query = function(batch\ view)$$

The speed layer takes care of the real-time streaming data. If we continue with our previous example of the daily batch view generation, then the speed layer will process today's data that was not in the scope of the previous batch run. Thus, the speed layer takes care of the data between two consecutive batch runs to generate the real-time view. The major difference between the batch layer and the speed layer is that the speed layer process most recent data on its arrival, while batch layer processes all of the stored data at once. The other feature of the real-time view is that it keeps on updating with the new data [14].

$$real-time\ view = function(real-time\ view,\ newdata)$$

The serving layer stores the latest batch views and serves the results of the query. The final result of the query is the merged view of the most recent batch and stream view. The query result can be represented as follows [14].

$$query = function(batch\ view,\ real-time\ view)$$

From the system's point of view, each layer represents a separate subsystem as shown in Figure 1. The data feed goes to both the batch and the speed layer. The master data set persists in the batch layer while serving, and speed layers serve the final query. In some cases, the serving layer can also keep the latest real-time views [12] but this will need some additional features inside serving layer that we discuss in Section 3.3.

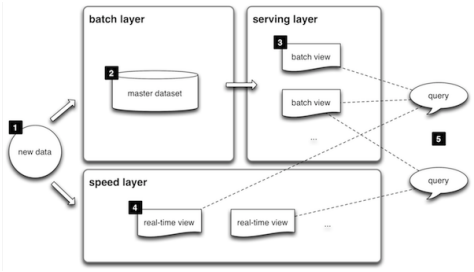

Figure 1: Lambda architecture high level overview [14]

2.1 Functional requirements of the real-time Big Data systems

Some characteristics such as scalability, fault tolerance, extensibility, maintainability, and interactivity,
etc. are associated with all kinds of real-time Big Data systems. The scalability refers to the ability of such systems to scale with the data volumes and other system requirements. The ability of the data systems to horizontally scale or scale out [1] provides a cost effective way of extending systems according to needs. Secondly, in distributed systems, it is also important to tolerate and recover automatically from some machines or subsystem failures. In case of such failures, systems should be able to ensure the completion of tasks as well as they should be able to auto-recover from such failures. Similarly, it is almost inevitable to avoid the human mistakes either in the deployment of systems or during the implementation of business logic. There should be some inbuilt mechanism in Big Data systems to rectify such mistakes.

The Big Data systems evolve with changes in the existing requirements as well as new features. In such cases, our target system should be flexible to extend easily [14]. Maintenance is the work required to keep the system running smoothly. The Big Data systems should be designed to minimize this overhead. Lambda architecture promotes the idea of keeping complexity out of the core components and put into the components with discardable output [14]. It makes debugging very easy.Interactivity refers to the response time of a query defines the interactivity of a particular data system. Various researches on human-computer interaction design recommend response time to be less than 10 seconds for any real-time interactive system [15].

[1] When the need for computing power increases, the tasks are divided among a large number of less powerful machines with (relatively) slow CPUs, moderate memory amounts, moderate hard disk counts.

3 Design trade-offs and implementation choices

Lambda architecture provides the basic architectural principles and guidelines for real-time Big Data systems. For building a real-time data system, there are many available tools that we can use in the subsystems or layers of the Lambda architecture. The design aspect of all the components should be considered to make the right architectural choices. The data analysis use cases drive the requirements. In our study, we have considered the requirement of each subsystem, evaluated different choices of components and provided recommendations along with a prototype implementation of lambda architecture while discussing the trade-offs.

3.1 Data collection and Ingestion

Data collection from multiple sources and ingestion into some persistent storage is typically the first step in building a Big Data system. The real-time streaming data requires the continuous collection of data. Such a data system should have inbuilt resilience against the delays and other imperfections such as missing, duplicated or out of order data [18]. It should also be able to collect data from the multiple sources and have the ability to channel collected data towards stream processing subsystems and data storages. In either case, data ingestion must ensure that it sends data at least once. Scalability in data collection and ingestion means adding more capacity, as well as the number of data sources. While extensibility in such systems refers to ability to handle multiple data formats and serialization/ De-serialization mechanisms.

In the Lambda architecture, collected data is sent to both the batch and the speed layers. The data on the batch layer resides in a distributed file system or distributed database. On the other hand, stream processing engine consumes the data on speed layer in real-time. Thus, the data collection and ingestion mechanism needs to guarantee the delivery of data on both channels.

Message brokers and data collection tools such as Apache Kafka, Apache Flume, Scribe, RabbitMQ, etc. can be used inside the Lambda architecture. All of these tools handles data in different ways, and the design aspects of these tools need consideration before using them in Lambda architecture. In our prototype implementation, we used Apache Kafka for data collection in combination with Apache Flume for data ingestion. Both of these tools are distributed, scalable, and fault tolerant.

Apache Kafka is a message broker that provides a pull-based model for data delivery [11]. The pull-based model helps applications to consume data at their pace. Apache Kafka also guarantees that data delivers at least once [11]. If the consumer system runs without any faults, then Kafka can also ensure the exact once delivery. However, if a process dies at the consumer end, then Kafka can send duplicate data to the process that takes over the tasks of the dead process. Apache Zookeeper [8] provides the coordination for running Kafka in distributed mode. The data replicates to multiple servers registered with Zookeeper for the purpose of fault tolerance. Zookeeper itself runs in a fault tolerant configuration and ensures high availability.

Using Apache Flume with the Apache Kafka, provides added advantage for data ingestion because of its ease of integration with Apache Hadoop (in the batch layer). It simplifies maintenance of the whole system. Although, one can argue that the addition of a system must increase the maintenance overheads. But since Apache Kafka does not provide an inbuilt mechanism for integration with Apache Hadoop, so we need to have an add-on for this purpose. Only using Apache Flume is another option but the Flume memory channel has weak protection against data loss in case of failures [7]. In this way, we can use the strengths of both systems while overcome the weaknesses.

3.2 Batch Layer

There are two main tasks of this layer (i) maintain the master copy of data and (ii) precompute the batch views.

Storing data requires a distributed file system or a distributed database that can ensure scalability on demand, fault tolerance, and easy maintainability. Hadoop filesystem (HDFS) is one of the most commonly used file system in Big Data landscape. HDFS is highly scalable, and it can scale up to thousands of nodes for parallel processing of data [16]. Reliability and fault tolerance is provided by keeping multiple copies of data on different servers. Tools such as Apache Hive, Pig, and Apache Spark can be used to extract, transform and load data on top of HDFS.

In Lambda architecture, each batch run includes the complete data set [14], which is a resource intensive task but it provides extra fault tolerance and a way to implement the new business logic as seamlessly as possible. We will discuss it in the Section 4. Lambda architecture also recommends creating a new view as latest batch view instead of updating the previous one because the updating requires random writes that are more expensive than the sequential writes in case of new view creation.

In our prototype implementation, we are using HDFS to keep the master copy of data. On top of HDFS, we are using Apache Hive for data warehousing and Apache Spark for data processing and batch view creations. Apache Hive runs on top of the Apache Hadoop and can load projections of data directly from HDFS in the form of tables similar to SQL-based database management systems. It provides a query language called Hive Query Language (HQL) [20] which is also very similar to SQL. Although HQL can be used to transform data by running MapReduce jobs in the background. In our setup, we are using Hive for housekeeping activities like partitioning and loading the data. We then use Apache Spark to process this partitioned data. Apache Spark is much faster than Apache Hive or MapReduce because it provides data abstraction that can run in-memory computations in a distributed parallel fashion [21]. This data abstraction is called the Resilient Distributed Datasets (RDD). Apache Spark also satisfies all the basic requirements of distributed big data system that we had mentioned earlier, and it comes with multiple programming language support such i.e. Scala, Java, and Python.

The batch view created as an output of the batch layer loads into to the serving layer for temporary storage till the

next most recent batch view is available.

3.3 Serving Layer

The main purpose of the serving layer is to provide the results of a query with minimum latency. As mentioned by Nathan Marz, for the case of serving layer the tooling lags behind the theory [14]. Thus, any generic distributed database system that can serve with low latency can be a candidate for serving layer. Use cases and CAP theorem [2] impacts the choice of tools. The serving layer tightly integrates with batch layer and contains the latest batch views. Due to the latency of batch execution the serving layer usually lags behind the real-time. The speed layer covers this lag.

In our implementation of Lambda architecture, we used HBase with OpenTSDB (open Time Series Database). HBase is an open-source implementation of the Google's BigTable [9]. It is a large-scale distributed columnar database [3]. It provides a very efficient table scan mechanism, which can serve queries with quick responses. To gain even faster responses, we have added OpenTSDB as a lightweight projection on top of HBase. OpenTSDB implements a smart way of storing time series data in HBase [17]. OpenTSDB simplifies HBase schema design and data maintenance through its very simple to use HTTP API. It is important to remember that OpenTSDB is best for time series data only.

3.4 Speed Layer

The purpose of the speed layer is to process real-time streaming data. The speed layer takes care of data for the time between two consecutive batches. The real-time processing requires incremental computing that adds some extra challenges as compared to batch and serving layers. The first and most obvious challenge is the requirement for frequent view updates. Then unlike batch layer real-time views need random writes for incremental changes that can have more latency as compared sequential writes. The scalability and fault tolerance requirements require distribution and replication of views across the cluster(s).

For use in Lambda architecture, there are various choices available as stream processing engines such as Apache Storm, Apache Spark Streaming, and Microsoft Trill, etc. In our study, we have compared Apache Storm with Apache Spark Streaming and selected Apache Storm for our prototype implementation. Although both of these systems have the capability to process real-time streams in distributed parallel way, Apache Spark Streaming differs from most of other stream processing engines by basic design ideology. It processes real-time stream in the form of micro batches using a data abstraction called discretized streams (D-Streams) [22]. Both streaming engines are capable of handling real-time streams on a sub-second level, provides hand full of nice transformation functions to process data. They both also provide support for multiple programming languages, with Storm capable of interfacing with more languages than Apache Spark e.g. Ruby, Nodejs, Clojure, etc.

We selected Apache Storm for use in our prototype setup because of better fault tolerance mechanisms than Apache Spark. In Apache Storm, if a node dies then the Nimbus service (similar to job tracker in Hadoop) reassigns the job to another worker node. The Nimbus and supervisor (similar to task tracker in Hadoop) services are designed to fail fast and recover automatically. Storm uses Apache Zookeeper that has the nimbus and supervisor states. When the Nimbus service is being restarted, the data processing job keeps on running. The only degradation in such case will be the unavailability of Nimbus for new tasks till it is running up again [19]. Apache Spark streaming also provide fast recovery for their driver (master node). In some reported cases, it was observed faster than storm [22] in this regard but the state of driver component can be lost during the failure, that results in loss of data during that brief time. Apache Spark has released a write-ahead log feature for Apache Spark Streaming with version 1.3 [5], However enabling this feature needs some compromises on data processing latency. This feature also needs further evaluation.

3.5 Prototype implementation

We have implemented a prototype real-time Big Data system as a proof of concept of Lambda Architecture. We used Apache Kafka for data collection. It was integrated with Apache Flume for data ingestion to Batch layer while a Kafka-Storm spout was configured to send data to speed layer. The batch layer was implemented using Apache HDFS for storing data, Apache Hive for data-warehousing and Apache Spark for batch processing. We used Apache HBase with OpenTSDB to store batch views in serving layer. The speed layered consisted of Apache Spark as a real-time data stream processing engine. The incremental real-time views from speed layer also resided inside Apache HBase. R project along with R-shinny was used to merge the views, and provide real-time data visualizations in the form an interactive dashboard. Figure 2 shows our prototype Lambda architecture implementation.

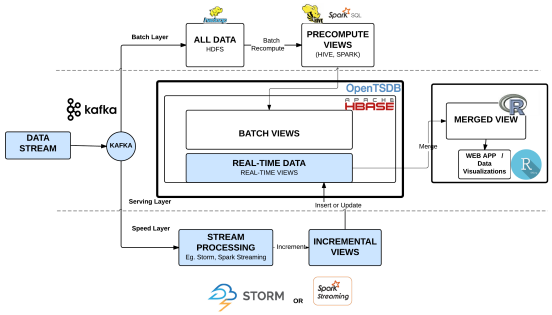

Figure 2: Lamdba architecture demo setup

We used Cloudera Hadoop (CDH-5.3.3) distribution as an integrated environment for the batch and serving layer com-

[2] CAP theorem states that, in the distributed computing environment it is impossible to guarantee consistency, availability, and partition tolerance simultaneously [3].

[3] a database that stores data tables as a section of columns instead of rows.

ponents. To take benefits from new DataFrame API and Apache Spark-SQL features we upgraded Apache Spark to version 1.3.1. For serving layer OpenTSDB was installed separately on top of HBase in CDH. Apache spark cluster was installed to run with supervisor [1]. We used OpenStack cloud infrastructure for the deployment of the complete prototype system. Table 1 and Table 2 summarize the prototype implementation software and infrastructure setups respectively.

Subsystem	Software Component	Version
Data Collection	Apache Kafka	0.8.1
Batch Layer	Apache Flume	1.5
	Apache Hive	0.13.10
	Apache Spark	1.3
	Apache Hadoop	2.5
Serving Layer	Apache Hbase	0.98.6
	OpenTSDB	v2.1RC
Speed Layer	Apache Storm	9.3

Table 1: Summary of softwares

Subsystem	Infrastructure
Data Collection, Batch Layer, Serving Layer	1 x Cluster with 1 x 4vCPU, 15GB RAM, 200GB hpc storage. 3x 1vCPU,3.5GB RAM, 200GB hpc storage
Speed Layer	1 x Cluster(1 x Nimbus + 3 x worker nodes) 4x 1vCPU,3.5GB RAM, 200GB hpc storage

Table 2: Summary of infrastructure

3.6 Data processing and work-flows

The data set, processed in our prototype setup was the traffic data collected by sensors installed in a highway. Each sensor provided data related to one lane of the highway. Out of various parameters provided by the sensors, we used the speed of the vehicle for our selected use case. As a demo use case, we were continuously calculating the rolling medians of the speed of vehicle per each lane in real-time. The rolling medians time windows were 1,5 and 10 seconds respectively. This particular use case was selected to test our system's capability to process data in parallel as well as identify the computation functions that can not execute in distributed parallel fashion.

3.6.1 Work-Flows

The continuous data processing pipeline consisted of following work-flows.

Data Collection: We collected the data using Apache Kafka. Within Kafka producer, we had implemented a web scraping bot that can pull the data from an HTML table. The frequency and the time window for data collection were configurable. It was important to tune our collection mechanism in a way that we do not miss any data. Replication of data is easy to handle in data processing steps. As mentioned in Section 3.1 the data was sent simultaneously to two channels i.e. batch layer and speed layers. This was done by configuring the Kafka producer to send data to two different topics.

Batch Layer Storage: On batch layer data is consumed using Apache Flume. Flume consumed the data from a given topic and ingested it directly into HDFS on a configured in a path. It was also configured to pull data from Apache Kafka as soon as new data is available. We used Apache Flume agent configurations to set up the size and rotation policy for the HDFS files.

Speed Layer Consumption: Kafka-Spout was used to integrate Apache Kafka with Apache Storm. Similar to Flume, Kafka Spout pulled the data from the Kafka as soon as data is available and provides it to the storm aggregate bolt.

Batch Layer Data Warehousing: Apache Hive loads data from HDFS into a partitioned table. The partitions were created on a daily basis and contained one day of data. This is extremely simple and efficient method to control what data to include while running the batch view. The crontab functionality of the Linux operating system was used to load of data in an automated way. As an improvement, the work-flow scheduling for Hadoop systems can be managed through Apache Oozie.

Batch views and serving layer: The data processing for our use case was done using PySpark (Python for Apache Spark). Apache spark can load data directly from Apache Hive tables and transform data as required. For this use case, we were utilizing spark to pre-process data and then we used Python Pandas library to calculate the rolling medians sequentially. The calculated values with the timestamps were sent to OpenTSDB as part of the batch view using HTTP API. OpenTSDB stores data in HBase and simplifies schema generation and database management.

Speed view and serving layer : Data processing on speed layer was done using an aggregate bolt. The data transformation logic was similar to batch, and we also used Python Pandas library for rolling mean calculation in streaming layer.

Merging the views and Visualizations: The merging of the batch and real-time views formulate the final result of the query. The batch jobs were run daily while real-time view took over for the data that had arrived after the start of the last batch. We implemented the merging logic in R programming language. R connects to OpenTSDB via HTTP API and pulls the recent versions of both the views. An interactive dashboard was created using R-shiny on top of R to provide an interactive environment for users. Dashboards showed the most recent rolling median speed time series graphs (one graph per highway lane). Other interactivities, such as rolling median time window selection, and time of days selections were also possible. Each interactivity generated a query for the most recent batch view and most recent real-time view OpenTSDB served the results in real-time and R merged the data and created the respective visualization.This interactive dashboard is illustrated in Figure 3.

4 Discussion

Lambda architecture provides a workable model for real-time Big Data system by combining the power of batch

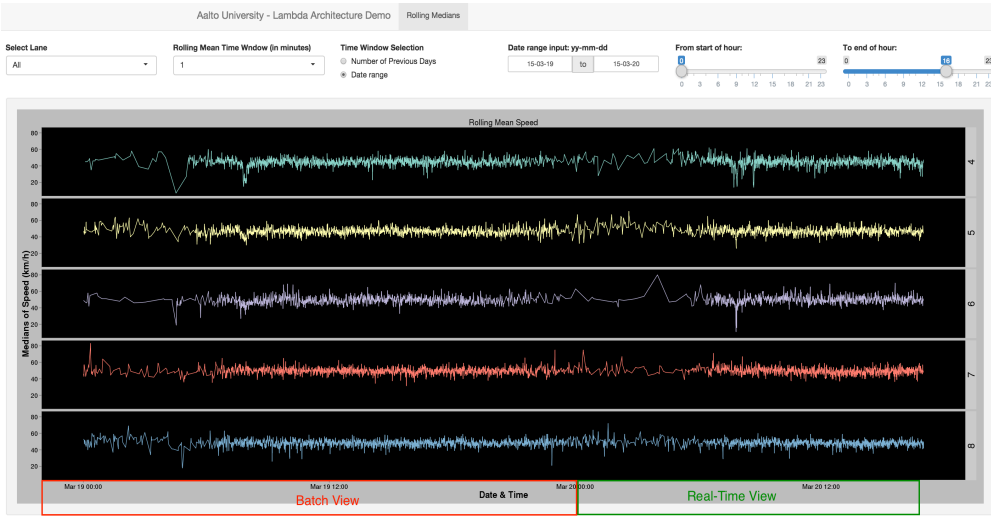

Figure 3: Live interactive dashboard

processing and speed of stream processing systems. It achieves the overall aim of low latencies while processing large amounts of data and provide real-time interactivity. The individual components of Lambda architecture can scale independently on demand basis. The architecture itself is very flexible and allows a variety of different available tools to fit in the framework. The data analysis use cases drive the requirements and thus the choice of tools within the different layers of Lambda architecture. As an example, we could have used Apache Cassandra instead of Apache HBase in our prototype implementation if our use case required availability and partition tolerance more than consistency.

Another very important and positive feature of lambda architecture is its inbuilt mechanism for recomputing data over and over again [10]. Firstly this allows easy changes inside business logic and secondly it provides an extra fault tolerance mechanism against human mistakes and coding bugs. With the re-computation of next batch view, the required changes and rectifications become effective. This feature of lambda architecture adds a lot to the extensibility of a data system. The master dataset also keep intact and immutability of data is ensured in lambda architecture [2].

With all the advantages of Lambda architecture, there are few weak points associated with it as well. By far the most highlighted inherent weakness of Lambda architecture is the requirement of maintaining the business logic in two different layers. That means coding the same use case twice for possibly two different system e.g. in our prototype; we had to code in PySpark for batch views and then re-code it in speed layer over Apache Storm. Using Apache Spark and Apache Spark streaming together provides an option to mitigate this weakness. However in Section 3.4, we have discussed the fault tolerance issue of Apache Spark Streaming. The ongoing work on Spark Streaming's fault tolerance seems very promising, and Apache Spark may be able to provide an enhanced unified environment in near future. This is also a new direction for further research on Lambda architecture and fault tolerance of streaming engines.

Another common critique of the Lambda architecture is its principle of recomputing batch views constantly [10]. As highlighted before, re-computing is good for fault tolerance and change management. The question arises, Do we need to recompute everything in every next batch run? Or Do we only recompute in for change in business logic? The counter argument in favor of Lambda architecture is that constant re-computation act as autoimmune for faults in speed layer.

Lambda architecture provides the basic construct of a real-time Big Data system. There is a flexibility to choose the components within the three defined layers. There are no specified limitations of using any pattern of the components. In light of the CAP theorem [4], on a very conceptual level, there may exist some patterns and anti-patterns of using various components inside lambda architecture. For example in the batch layer and serving layer, we typically emphasize on consistency over availability and vice versa in speed layer. So in case we use a single AP (available and network partition tolerant) component in batch or serving layer than we may lose consistency for the complete batch view. Similarly, if we use a consistent component in speed layer than we may lose real-time in real-time view. Mapping this concept on our demo setup, we are using Kafka in AP configuration, that may compromise the consistency of our batch views. Lambda architecture neither dictates nor specify any of such theoretical principles.

It is worthwhile here to mention that there are some emerging architectures that tend to solve the problems of Lambda architecture. One of them is the Kappa Architecture, proposed by Jay Kerps, who is also one of the main contributors to Apache Kafka project. He proposes a powerful all stream processing engine instead of two-layered architecture [14].

References

[1] Supervisor: A process control system.

[2] The lambda architecture: principles for architecting realtime big data systems, 2013.

[3] E. Brewer. A certain freedom: thoughts on the cap theorem. In *Proceedings of the 29th ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, pages 335–335. ACM, 2010.

[4] E. A. Brewer. Towards robust distributed systems. In *PODC*, volume 7, 2000.

[5] T. Das. Improved fault-tolerance and zero data loss in spark streaming, January 2015.

[6] B. Ellis. *Real-time Analytics: Techniques to Analyze and Visualize Streaming Data*. John Wiley & Sons, 2014.

[7] J. Gwen Shapira. Flafka: Apache flume meets apache kafka for event processing.

[8] P. Hunt, M. Konar, F. P. Junqueira, and B. Reed. Zookeeper: Wait-free coordination for internet-scale systems. In *USENIX Annual Technical Conference*, volume 8, page 9, 2010.

[9] A. Khetrapal and V. Ganesh. Hbase and hypertable for large scale distributed storage systems. *Dept. of Computer Science, Purdue University*, 2006.

[10] J. Kreps. Questioning the lambda architecture, 2014.

[11] J. Kreps, N. Narkhede, J. Rao, et al. Kafka: A distributed messaging system for log processing. In *Proceedings of 6th International Workshop on Networking Meets Databases (NetDB), Athens, Greece*, 2011.

[12] MAPR. Lamnda architecture: Making sense of it all.

[13] N. Marz. How to beat the CAP theorem, 2011.

[14] N. Marz and J. Warren. *Big Data: Principles and best practices of scalable realtime data systems*. Manning Publications Co., 2015.

[15] J. Nielsen. *Usability engineering*. Elsevier, 1994.

[16] K. Shvachko, H. Kuang, S. Radia, and R. Chansler. The hadoop distributed file system. In *Mass Storage Systems and Technologies (MSST), 2010 IEEE 26th Symposium on*, pages 1–10. IEEE, 2010.

[17] B. Sigoure. Opentsdb: The distributed, scalable time series database. *Proc. OSCON*, 11, 2010.

[18] M. Stonebraker, U. Çetintemel, and S. Zdonik. The 8 requirements of real-time stream processing. *SIGMOD Rec.*, 34(4):42–47, Dec. 2005.

[19] A. S. Team. Apache storm: Fault tolerance, 2015.

[20] A. Thusoo, J. S. Sarma, N. Jain, Z. Shao, P. Chakka, S. Anthony, H. Liu, P. Wyckoff, and R. Murthy. Hive: A warehousing solution over a map-reduce framework. *Proc. VLDB Endow.*, 2(2):1626–1629, Aug. 2009.

[21] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica. Spark: cluster computing with working sets. In *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing*, pages 10–10, 2010.

[22] M. Zaharia, T. Das, H. Li, T. Hunter, S. Shenker, and I. Stoica. Discretized streams: Fault-tolerant streaming computation at scale. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*, pages 423–438. ACM, 2013.

Cloud datastores: NewSQL solutions

Dmytro Arbuzin
Student number: 415158
dmytro.arbuzin@aalto.fi

Abstract

Distributed data stores have undergone a huge change over the last decade: from the first NoSQL databases to modern scalable cloud systems, which serve the major part of today's Internet data. The release of Google's BigTable paper in 2006 has inspired many developers and companies to move towards NoSQL solutions. That generated a huge push in developing distributed data stores and significantly improved their performance and scalability. However, the industry could not overcome the limitations of distributed databases, specifically the lack of ACID semantics. During the last few years many so-called NewSQL databases were presented which combine scalability and performance of NoSQL data stores with sustainability and reliability of traditional databases. This paper evaluates the number of existing ACID-compliant cloud datastores and discusses the main difficulties and challenges in implementing ACID transactions in distributed data stores.

KEYWORDS: Database, NewSQL, ACID, transaction, NoSQL, CAP

1 Introduction

The database industry used to be the most stable and steady in whole IT world. It used to be the case until the beginning of 2000s, when some data sets could no longer be stored on a single computer. That caused a huge shift from traditional relational SQL database management systems towards highly scalable distributed systems. New design solutions were required due to a number of limitations that distributed systems possesses. Those limitations were initially formulated by Erik Brewer back in 2000 and now are widely known as CAP theorem [10].

In 2004 Google started to deploy their own distributed data storage to maintain most of internal and external services. Later, in 2006, they published their BigTable paper [13], finally introducing a first successful fully distributed, scalable data storage to the external world. One year later, Amazon presented their distributed highly available key-value storage called Dynamo [17]. Now, many experts state that BigTable has inspired the whole industry to actively move towards NoSQL solutions. Rick Catell in his work [12] wrote that BigTable, Dynamo and Memcached, scalable distributed memory caching system, had provided a "proof of concept" of NoSQL solutions and motivated many data stores that exist today. For example, an open-source, distributed Apache HBase system, which is currently a platform for a Facebook messaging service [20], was a direct clone of BigTable.

However, even being a huge success, BigTable could not provide enough features that are required for modern web services. Current Internet applications, among other components, require *high scalability*: the ability to grow fast while serving millions of users; *low latency*: the response time of website should be as low as possible; *consistent view of data*: users should be able to read their writes immediately; and finally the application should be *highly available*: websites should run 24/7 despite all potential problems. However, all these features together create a conflict: most of them can be easily achieved with traditional centralized SQL databases, but such databases can not scale horizontally well to a large number of users. From the other side, NoSQL datastores like BigTable or HBase can not ensure strong consistency, low latency and other critical features at the same time. In other words, NoSQL databases do not follow ACID semantics, like relational ones, but they support horizontal scaling or scaling out design principle. Scaling out basically means the ability to process over many hardware nodes.

Google tried to resolve this conflict and launched Megastore [6] in 2006: the database system build on top of BigTable with support of full ACID semantics within partitions, but only limited consistency guarantees across them. The idea was that one partition inside cloud storage should be enough for most Internet applications. Megastore had made the life of developers easier, even though it did not solve the conflict completely. It was complex and comparatively slow for operations across partitions.

In 2012, Google presented a key/value store called Spanner [15], which used to be the first distributed database with consistent transactions across nodes. One year earlier in 2011 the NewSQL term was introduced by Matthew Aslett [3].

This paper discusses the main design challenges of NewSQL datastores and describes the architecture of some existing solutions.

The remainder of the paper is structured as follows: Section 2 explains core concepts of database architecture, such as CAP theorem and ACID principles. It also introduces vocabulary used in the paper. Section 3 discusses the main design challenges and trade-offs in achieving ACID semantics in NoSQL databases. Section 4 discusses Spanner and F1 systems from Google. Section 5 discusses CockroachDB: an new open-source NewSQL solution. Section 6 summarises the paper.

2 Main concepts of database architecture

This section introduces and explains some of the most important terms and core concepts of database design.

2.1 Terminology

Because of the huge variety of products and technologies there is often a terminology collision, which leads to misunderstanding ans confusion. It is not a secret that many companies use "trendy" words such as "cloud", "NoSQL", "distributed", etc. for marketing purposes. There is no common terminology in academics either. Thus, it is necessary to clarify some basic vocabulary, which is going to be widely used in the following discussion.

Traditional database (RDBMS) is a relational SQL database system which follows ACID requirements.

In his article [12] Rick Cattell has described NoSQL or "Not only SQL" [18] with six key features. Based on that description, we can define NoSQL as a scalable distributed database system, designed mainly for OLTP (Online Transaction Processing) handling using concurrency model weaker than ACID transactions. This definition is quite broad and might be arguable, but it satisfies us in the context of this paper.

Based on the previous two definitions, we can describe NewSQL as a class of databases which provide the same scalability and throughput performance as NoSQL systems while still maintaining the ACID guarantees of a traditional database systems.

2.2 ACID

The current concept of transaction was formulated in the 1970th, in the era of early database development. We can define it as a set of operations grouped in one working unit. Later Theo Haerder completed transaction's description using the acronym ACID [19], which is one of the most important concepts in database theory. It stands for atomicity, consistency, isolation and durability. Each database must strive to achieve all four of these features, only then the database can be considered reliable [14].

- Atomicity. The transaction should follow the "all-or-nothing" rule. In other words, all changes must be stored to the system, or no changes at all in case of any error or conflict during the transaction execution.
- Consistency. Each successful transaction by definition commits only legal results. The data view state is the same for all database connections.
- Isolation. Events within a transaction should not impact other transactions running concurrently. As Concurrency is a key feature of distributed systems, isolation is a point where the most trade-offs are made. See Section 3.3.
- Durability. Once the transaction has been executed, the results will be stored permanently.

Almost all modern SQL databases have ACID-compliant transactions. They provide the essential level of reliability for most important industry areas, such as banking systems for example.

2.3 CAP theorem

NoSQL systems generally strive for scalability and performance, and as a result they do not match ACID requirements. This phenomenon was firstly formulated by Eric Brewer back in 2000 and today it is widely known as a CAP theorem [10]. It states that in a distributed system, you can only have two out of the following three guarantees across a write/read pair: Consistency, Availability, and Partition Tolerance - one of them must be sacrificed.

- Consistency. A read is guaranteed to return the most recent write for a given client. All nodes have the same state of data view. It is important to understand that CAP consistency is not the same that ACID consistency.
- Availability. A non-failing node will return a reasonable response within a reasonable amount of time (no error or timeout).
- Partition tolerance. The system will continue to function whenever the node fails.

For distributed systems the acronym BASE is widely used in contrast to ACID. BASE stands for Basically available, Soft state, eventually consistent.

All distributed database systems are willing to ensure all three characteristics of the CAP theorem. By definition they ensure partition tolerance and it seems that the choice is straightforward: consistency or availability. However, with the development of NoSQL, trade-offs have became more complex. Currently, many distributed databases claim themselves as eventually consistent and/or highly available. The initial CAP principle "2 out of 3" became no longer accurate. In 2012, Eric Brewer explained his CAP theorem [9] and stated that the "2 out of 3" principle was misleading from the very beginning. Brewer also showed that by using the latest techniques, design approaches and flexible settings NoSQL systems can achieve a certain level of sustainability regarding all three principles. In other words, the CAP theorem still applies to distributed systems. However, with the usage of latest design techniques it is possible to achieve ACID semantics by a sophisticated balancing between Availability and Consistency. First of all, because both Consistency and Availability are not monotonic, but rather aggregated options consisting of many settings and rules and by relaxing some of them in Availability, the system can achieve stronger Consistency and vice versa.

3 ACID transactions in NoSQL: main difficulties

As discussed in previous parts of the paper, it is possible to achieve ACID transactions in distributed systems by using certain design techniques. However, also was mentioned that

the architecture of such systems is quite advanced and trade-offs are complex. This section discusses the main challenges, and also the compromises that are made at architecture level to achieve ACID semantics in distributed systems.

3.1 Latency

The main point of ACID semantics is to ensure the same behaviour and the same guarantees for distributed systems as for single node databases. However there is fundamental obstacle such as latency between nodes. Latency is a physical parameter which primary depends on the speed of light and can not be neglected by definition. In the best case, two servers on the opposite sides of the Earth, connected via cable will have a round-trip time 133.7 ms. Moreover, in real deployments of distributed systems, there are many secondary parameters such as routing, congestion, server-side overheads, etc, which increase the communication time between servers. Peter Bailis in this blog [4] demonstrates the average RTT between Amazon EC2 instances. The minimum is 22.5 ms and maximum is 363 ms. Taking into consideration that any operation execution over the partition also require a certain time, the overall transaction time can easily reach 1 second or even more.

All ACID options are interconnected among themselves and affects each other, however each of them also requires some specific set of actions to be achieved. Atomicity depends on a latency more than other options. For instance, if a transaction executes over multiple nodes, the overall speed of it depends on the slowest node and the system can not make the decision to commit or to rollback, until all nodes have responded. This might lead to the state of data being in an unsustainable mode for a relatively long time and ruin consistency, which ruins durability in the end. Appealing to CAP theorem, we can state that the main trade-off is not even Consistency versus Availability, but rather Consistency versus Response Time (RT). The more nodes a transaction involves, the higher the Response Time. RT affects Availability in a direct ratio, because to ensure a consistent view of a data system needs to use various concurrency control mechanisms and methods.

3.2 Concurrency control

Concurrency control is the activity of coordinating the actions of processes that operate in parallel, access shared data, and therefore potentially interfere with each other. In database theory, concurrency control ensures that database transactions are performed concurrently without violating the data integrity of the databases. Thus concurrency control is an essential element for correctness in any system where two or more database transactions, executed with a time overlap, can access the same data, e.g., virtually in any general-purpose database system.

Basically concurrency control ensures ACID requirements by certain mechanisms. This section describes the most important principles, methods and concepts that are used to achieve ACID not only on single-node servers but in distributed systems as well.

We can divide all concurrency mechanisms into three main categories:

- Optimistic - the family of methods, which assume that transaction conflicts are highly unlikely. It does not block (lock) any data during transaction execution and checks whether a transaction meets the isolation and other integrity rules in the end of its execution [23]. If there were any rules violations transaction rollbacks and re-executes, otherwise it commits. As many current systems prefer Performance and Availability over Consistency, most of NoSQL databases use optimistic approach. However, if conflicts are frequent and rollbacks occur often it is not sufficient to use optimistic methods.
- Pessimistic - blocks entire data units, which are involved into transaction, until it ends. Blocking operations typically reduces performance and also can lead to deadlocks.
- Semi-optimistic - block operations in some situations, if they may cause violation of some rules, and do not block in other situations while delaying rules checking to transaction's end.

The crucial difference between these mechanisms is performance, which consists of many factors, such as average transaction completion rates (throughput), transaction types mix, computing level of parallelism and others.

The mutual blocking between two transactions or more results in a deadlock, where the transactions involved are stalled and cannot reach completion. Most non-optimistic mechanisms are prone to deadlocks which are resolved by an intentional abort of a stalled transaction (which releases the other transactions in that deadlock), and its immediate restart and re-execution. The likelihood of a deadlock is typically low. Blocking, deadlocks, and aborts all result in performance reduction, and hence the trade-offs between the categories [30].

To achieve consistent view of a data the concurrency control has many methods and algorithms. The most popular is Two-phase locking. However for distributed systems often more advanced algorithms are required and even combinations of several.

Two-phase locking (2PL) is a traditional algorithm that guarantees Serializability - the highest level of data isolation. The algorithm consists of two phases: in the first expanding phase it locks all data units one by one. In the second phase it releases the locks. Figure 1 illustrates the main principle of this algorithm. 2PL is pessimistic and reduces performance heavily. In it's original state it is also not protected against deadlocks.

Multiversion concurrency control (MVCC) is a concept that can guarantee Snapshot Isolation - the highest level of data isolation in distributed systems. Being an optimistic MVCC does not lock data, but provides a different data view to overlapping transactions by taking snapshots and marking them with timestamps. In this way, there is a possibility to store different versions of the same data unit.

The main aim of concurrency control algorithms is to ensure certain level of data Isolation in terms of ACID terminology.

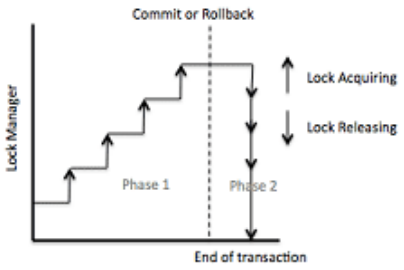

Figure 1: Two-phase locking algorithm [21]

3.3 Isolation

Isolation is probably the most important ACID option in terms of distributed transactions. The reason behind this is that Atomicity, Consistency and Durability are well-defined values and always remain to ensure the same result. In contrast to them, Isolation level can be different. According to the ISO standard, there are 4 levels of isolation [7]. Furthermore, there are few more levels, which are not described by the standard, but are commonly used in industry. Table 1 explains different levels according to read phenomena, which can occur on each level. These phenomena are beyond this paper's scope, however they are widely known in terms of database theory and defined by ISO SQL-92 standard.

Isolation Level	**Dirty Read**	**Fuzzy Read**	**Phantom**
Read uncommited	Possible	Possible	Possible
Read commited	Impossible	Possible	Possible
Cursor stability	*Impossible*	*Sometimes possible*	*Possible*
Repeatable read	Impossible	Impossible	Possible
Snapshot isolation	*Impossible*	*Impossible*	*Sometimes possibe*
Serializable	Impossible	Impossible	Impossible

Table 1: Levels of isolation in terms of three original phenomena. *Italic* - not defined by the ISO standard

Table 1 shows that Serializable level is the most stable and ensures the highest safety of the data. This level means that transactions can not overlap between each other and usually uses two-phase locking mechanism [8]. The database designers have realized long time ago that Serializability can not be achieved in distributed systems [16]. However, Table 2 shows that many current industrial RDBMS do not provide Serializability by default or even do not have it at all. This is because such isolation level limits concurrency options tremendously even on a single node, which is unacceptable for modern applications.

Information from Table 2 is also particularly surprising, when we consider the widespread deployment of many of these non-serializable databases, like Oracle 11g, which are well-known to power major business applications. Taking into consideration that most of RDBMS vendors have already agreed to relax Isolation level to achieve better concurrency performance, we can state that inability to achieve

Database	**Default**	**Max**
Actian Ingres 10.0/10S	S	S
Aerospike	RC	RC
Akiban Persistit	SI	SI
Clustrix CLX 4100	RR	RR
Greenplum 4.1	RC	S
IBM DB2 10 for z/OS	CS	S
IBM Informix 11.50	Depends	S
MySQL 5.6	RR	S
MemSQL 1b	RC	RC
MS SQL Server 2012	RC	S
NuoDB	CR	CR
Oracle 11g	RC	SI
Oracle Berkeley DB	S	S
Oracle Berkeley DB JE	RR	S
Postgres 9.2.2	RC	S
SAP HANA	RC	SI
ScaleDB 1.02	RC	RC
VoltDB	S	S
RC: read committed, RR: repeatable read, SI: snapshot isolation,S: serializability, CS: cursor stability, CR: consistent read		

Table 2: Default and maximum isolation levels for ACID DBMS as of January 2013 [5]

Serializability in NoSQL systems appears as non-critical for practical applications.

Table 1 also presents that the second most reliable Isolation level is SI (Snapshot Isolation). A transaction executing with Snapshot Isolation always reads data from a snapshot of the (committed) data as of the time the transaction started, called its *Start-Timestamp* [7]. When the transaction concludes, it will successfully commit only if the values updated by the transaction have not been changed externally since the snapshot was taken. Such a write-write conflict will cause the transaction to abort. Even though SI provides strong data consistency, some anomalies such as *skew write* may still occur [7]. In 2009, Michael Cahill showed [11] that skew write anomalies in SI level could be prevented by detecting and aborting "dangerous" triplets of concurrent transactions. New level is known as Serializable Snapshot Isolation (SSI). SI and SSI are implemented within MVCC as usual.

3.4 Consensus Protocols

The major part of distributed systems theory are consensual protocols. Two most popular protocols are Paxos and Two-phase commit (2PC).

Tho-phase-commit protocol

2PC protocol is one of the most popular and well known atomic commitment protocol in distributed systems. It was introduced already in 1970th and was widely utilized by industry due to it's relatively simplicity and cheap cost in terms of number of operations [32]. This protocol can effectively solve the number of possible problems and conflicts during transaction executions including process, network node, communication and other failures. Many protocol variants exist. They primarily differ in logging strategies and recov-

ery mechanisms.

The basic protocol operates in following manner: the system has one primary node, which designated as the coordinator; all other nodes operate as slaves or cohorts. Whole process is divided into two phases: first one is a *commit request phase* or a *voting phase*, during which the coordinator sends query to commit message to all nodes and waits for their answers. All nodes execute a transaction up to the moment when they need to commit and replies with the agreement message whether the execution was successful or not (basically they vote to commit or rollback the transaction). Also each node is processing logging for solving all possible problems. Second phase is *commit phase* or *completion phase*. If during voting phase all nodes answered positive the coordinator sends the commit message to all cohorts. Each cohort completes the operation, releases all locks and resources held during the transaction and sends the acknowledgment back. The coordinator completes the transaction when all acknowledgment are gathered. If during the voting phase even one node has answered NO in agreement message, the coordinator sends rollback message to cohorts, they rollback all operations using logs and send the acknowledgement back. The coordinator rollback the transaction after all acknowledgements have been gathered. This protocol has a couple of critical disadvantages. First of all it is a blocking protocol, which means that it reduces at least a write performance during the voting phase. Also with this protocol whole system is as fast as the slowest node. But the main problem that 2PC has a single point of failure, which is the coordinator. In case, the coordinator fails permanently, some cohorts will never resolve their transaction. There are variants of 2PC (dynamic two-phase commit for example), that are solving the coordinator problem, however that brings extra complexity and slows the algorithm.

Paxos

Paxos is a popular consensus protocol firstly proposed by Leslie Lamport [25], which is provably correct in asynchronous networks that eventually become synchronous, does not block if a majority of participants are available (withstands n/2 faults) and has provably minimal message delays in the best case.

Nodes can be in three roles: proposers, acceptors and leaders. The proposer sends 'prepare' requests to the acceptors. When the acceptors have indicated their agreement to accept the proposal, the proposer sends a commit request to the acceptors. Finally, the acceptors reply to the proposer noticing the success or failure of the commit request. Once enough acceptors have committed the value and informed the proposer, the protocol terminates. Many nodes can act as a proposers at the same time. Each proposal is unique and has a sequence number. In this way nodes can order the proposals by sequence number (time) and accept only new proposals. In practical Paxos implementations one of the nodes is elected as a leader. Leader is responsible for making progress. In this way system operates asynchronously. Paxos does not require all nodes to vote positively, but only the majority. Nearly half the nodes can fail to reply and the protocol will still continue correctly. Such effect is possible due to the fact that any two majority sets of acceptors will have at least one acceptor in common. Therefore if two proposals are agreed by a majority, there must be at least one acceptor that agreed to both. This means that when another proposal is made, a third majority is guaranteed to include either the acceptor that saw both previous proposals or two acceptors that saw one each. This method solves the blocking problem of 2PC. Also Paxos does not have single point of failure as if the leader fails the system can select a new one.

There are many versions and different implementations of Paxos exist. However, even fast ones are slower than 2PC, which is basically the fastest protocol. Also Paxos sacrifices liveness, i.e. guaranteed termination, when the network is behaving asynchronously, however this is only a theoretical problem.

Raft

Raft is a relatively new consensus protocol [29], which was developed as an alternative to Paxos. The primary reason for it was that Paxos is too complicated to understand and it does not solve efficiently current needs of distributed system, because it was designed in 1989 as a solution for mostly theoretical problems. In general, Raft is a successor of Paxos and follows the same logic, the main difference is in a leader election process. Raft requires leader election to occur strictly before any new values can be appended to the log, whereas a Paxos implementation would make the election process an implicit outcome of reaching agreement.

4 Google F1

Google F1 was created as a challenge to serve company's the most important direction - AdWords business. The F1 is a database system, which is built on top of Google key/value store called Spanner.

4.1 Spanner

Spanner is a multi-version, globally distributed and synchronously-replicated database [15]. In Spanner, inside one datacenter data is organized in a special way: rows are partitioned into clusters called directories using ancestry relationships in the schema. Directories contain fragments of data. Fragments of one directory are stored by groups. Each group has at least one replica tablet per a datacenter. Data is replicated synchronously using the Paxos algorithm and all tablets for a group store the same data. One replica tablet is elected as the Paxos leader for the group, and that leader is the entry point for all transactional activity for the group. There can be also readonly replicas, but they do not participate in leader election and cannot be leaders.

Spanner provides serializable pessimistic transactions using strict two-phase locking mechanism inside one group. Across multiple groups Paxos provides transactions using a two-phase commit protocol on top of Paxos. 2PC adds an extra network round trip so it usually doubles observed commit latency. 2PC scales well up to 10s of participants, but abort frequency and latency increase significantly with 100s of participants [31]. Google's philosophy behind this was stated in Spanner paper: *"We believe it is better to have application programmers deal with performance problems due*

to overuse of transactions as bottlenecks arise, rather than always coding around the lack of transactions" [15].

In general, Spanner has a set of interesting features, some of which were introduces by the first time ever. Firstly, the replication configurations for data can be dynamically controlled at a fine grain by applications, which allows user to control durability, availability levels and read performance by setting up number of replicas and their location. Secondly, spanner was the first system to provide externally consistent reads and writes, and globally-consistent reads across the database at a timestamp (Snapshot Isolation). The timestamps reflect serialization order. In addition, the serialization order satisfies external consistency (or equivalently, linearizability): if a transaction T1 commits before another transaction T2 starts, then T1's commit timestamp is smaller than T2's. On a large scale such guarantees are possible due to TrueTime API.

TrueTime API

For distributed systems using timestamps for synchronization time is a critical issue. Researchers introduced several solutions such as Logical Clock [24] and Vector Clock [27]. However, in practice most of current internet applications are using Physical Time and Network Time Protocol (NTP) for synchronization [28]. Since all of above systems are not suitable to maintain consistent distributed transactions for different reasons, which are beyond this paper's scope, Google decided to introduce their own solution.

TrueTime (TT) is system for time synchronization which consists of a complex mix of software and hardware. TrueTime uses a set of *time master machines* at each datacenter. The majority of masters uses GPS clock. The remaining masters (*Armageddon masters*) are equipped with atomic clocks. This is done because GPS and Atomic clocks have different failure mode. All masters' time references are regularly compared against each other. Each master also cross-checks the rate at which its reference advances time against its own local clock, and evicts itself if there is substantial divergence. Between synchronizations, Armageddon masters advertise a slowly increasing time uncertainty that is derived from conservatively applied worst-case clock drift. GPS masters advertise uncertainty that is typically close to zero [15].

Every daemon polls a number of masters to reduce vulnerability to errors from any one master. Both types of masters (GPS and Armageddons) are picked from different locations. Daemons apply a variant of Marzullo's algorithm [26] to detect and reject liars, and synchronize the local machine clocks to the nonliars.

Such complex architecture allows the system to reduce time uncertainly to 7 ms, which is enough to provide concurrency guarantees for external transactions.

4.2 F1 key features

Besides scalability, synchronous replication, strong consistency and ordering properties that F1 inherits from Spanner, F1 itself adds additional properties [31] such as:

- distributed SQL queries, including joining data from external data sources;
- transactionally consistent secondary indexes;
- asynchronous schema changes including database reorganizations;
- optimistic transactions;
- automatic change history recording and publishing.

Figure 2 represents basic architecture of F1 system. F1 is built on top of Spanner and F1 servers are located in the same datacenters to reduce latency. The Spanner servers retrieve their data from the Colossus File System (CFS) in the same datacenter. However, F1 can also communicate to Spanner servers in other datacenters to ensure availability and load balancing. F1 masters monitor the health state of a slave pool and distributes the list of available slaves to F1 servers. Operations from all clients go through F1 servers except MapReduce processes that executes directly on Spanner level to increase performance.

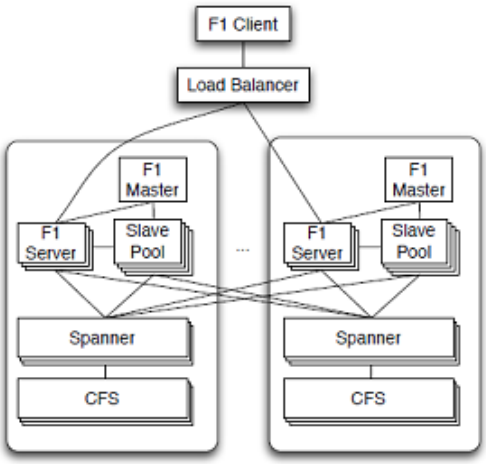

Figure 2: The basic architecture of the F1 system, with servers in two datacenters [31]

F1 provides three types of ACID transactions, all based on Spanner transactional support. Typical F1 transaction consists of number of reads and optionally one write to commit the result.

Snapshot transactions are read-only transactions, which are possible due to Spanner timestamp feature. Snapshot transactions allow multiple client servers to see consistent views of the entire database at the same timestamp.

Pessimistic transactions map directly on Spanner transactions.

Optimistic transactions consist of read non-locking phase and then short write phase. To solve conflicts F1 stores the modification timestamps for each row, in a special hidden locked column. In the end F1 server checks all row timestamps in a short-lived pessimistic transaction and send data to Spanner to commit the result if no conflicts occurs.

Optimistic transactions bring a lot of great benefits and performance to the system, however there are some trade-offs which come with them. Phantoms writes are possible and it can lead to a low throughput under high contention.

When many clients are incrementing one counter concurrently for example.

In the end, Spanner together with F1 brings database development on a new level and erase the boundaries between NoSQL and relational systems more than ever before.

5 CockroachDB

CockroachDB is a new scalable, transactional, geo-replicated datastore [1]. It is an open-source project leading by the team of "ex-Googlers", who were originally inspired by Google success and are willing to create database as powerful as Spanner, but not so complex to deploy and maintain. The project currently is on alpha stage and all documentation and progress can be found in their GitHub repository [2].

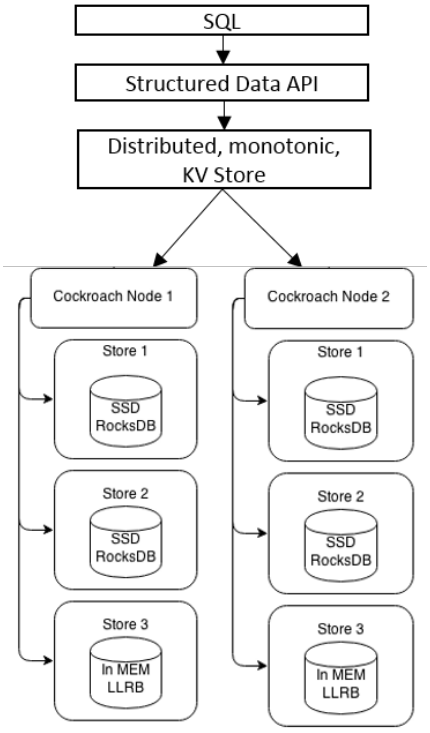

Figure 3: The basic architecture of CockroachDB consisting of 2 nodes

Cockroach is a distributed key/value datastore. Figure 3 represents the basic architecture of the system. It has layered structure. On the bottom level the system uses another open-source project called RocksDB to store actual data. The raw data is stored within sorted by key ranges inside stores. Ranges are replicated using Raft consensus protocol. One range has at least 3 replicas. One store is localed on one physical device. Sets of stores forms nodes. The distributed key/value store communicates with any number of physical cockroach nodes, forming the first level of abstraction by handling the details of range addressing. Then the the structured data API provides familiar relational concepts such as schemas, tables, columns, and indexes. SQL level is not implemented yet.

Cockroach provides optimistic distributed transactions. It has two levels of isolation: SI and SSI. Cockroach uses MVCC together with 2PC protocol to provide strong consistency. Each Cockroach transaction is assigned a random priority and a "candidate timestamp" at start. The candidate timestamp is the provisional timestamp at which the transaction will commit, and is chosen as the current clock time of the node coordinating the transaction. This means that a transaction without conflicts will usually commit with a timestamp that, in absolute time, precedes the actual work done by that transaction. In the course of organizing the transaction between one or more distributed nodes, the candidate timestamp may be increased, but will never be decreased. The core difference between the two isolation levels SI and SSI is that the former allows its commit timestamp to increase and the latter does not. Timestamps are generated by combination of both a physical and a logical components [22], which allows to provide low latency without using atomic and GPS clocks.

6 Conclusion

In this paper, we have discussed some core principles of the database theory and demonstrated some core design challenges of the transaction's implementation. Fast ACID-compliant transactions in distributed systems used to be an unachievable dream even 10 years ago. However, using the latest engineering solutions Google was able to implement ACID transactions on a global scale and provided the proof of concept to the whole world. Currently more and more systems strive to achieve NewSQL features. CockroachDB is a good example of such system. Even though, Cockroach is still under development and it is far from final realise, it provides a strong proof of concept that NewSQL datastores are possible to implement without additional hardware setups, simply by using the latest algorithms and techniques like HybridTime and Raft protocol.

References

[1] Cockroachdb. `http://cockroachdb.org/`. [Online; accessed: 2015-04-11].

[2] Cockroachdb github. `https://github.com/cockroachdb/cockroach`. [Online; accessed: 2015-04-13].

[3] M. Aslett. How will the database incumbents respond to NoSQL and NewSQL. *San Francisco, The*, 451:1–5, 2011.

[4] P. Bailis. Communication costs in real-world networks. `http://www.bailis.org/blog/communication-costs-in-real-world-networks`. [Online; accessed: 2015-03-10].

[5] P. Bailis, A. Fekete, A. Ghodsi, J. M. Hellerstein, and I. Stoica. HAT, not CAP: towards highly available transactions. In *Proceedings of the 14th USENIX conference on Hot Topics in Operating Systems*, pages 24–24. USENIX Association, 2013.

[6] J. Baker, C. Bond, J. C. Corbett, J. Furman, A. Khorlin, J. Larson, J.-M. Leon, Y. Li, A. Lloyd, and V. Yushprakh. Megastore: Providing scalable, highly available storage for interactive services. In *Proceedings of the Conference on Innovative Data system Research (CIDR)*, pages 223–234, 2011.

[7] H. Berenson, P. Bernstein, J. Gray, J. Melton, E. O'Neil, and P. O'Neil. A critique of ANSI SQL isolation levels. In *ACM SIGMOD Record*, volume 24, pages 1–10. ACM, 1995.

[8] P. A. Bernstein, V. Hadzilacos, and N. Goodman. *Concurrency control and recovery in database systems*, volume 370. Addison-wesley New York, 1987.

[9] E. Brewer. CAP twelve years later: How the" rules" have changed. *Computer*, 45(2):23–29, 2012.

[10] E. A. Brewer. Towards robust distributed systems. In *PODC*, volume 7, 2000.

[11] M. J. Cahill, U. Röhm, and A. D. Fekete. Serializable isolation for snapshot databases. *ACM Transactions on Database Systems (TODS)*, 34(4):20, 2009.

[12] R. Cattell. Scalable SQL and NoSQL data stores. *ACM SIGMOD Record*, 39(4):12–27, 2011.

[13] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber. Bigtable: A distributed storage system for structured data. *ACM Transactions on Computer Systems (TOCS)*, 26(2):4, 2008.

[14] M. Chapple. The ACID model. `http://databases.about.com/od/specificproducts/a/acid.htm`. [Online; accessed: 2015-02-05].

[15] J. C. Corbett, J. Dean, M. Epstein, A. Fikes, C. Frost, J. J. Furman, S. Ghemawat, A. Gubarev, C. Heiser, P. Hochschild, et al. Spanner: GoogleâĂŹs globally distributed database. *ACM Transactions on Computer Systems (TOCS)*, 31(3):8, 2013.

[16] S. B. Davidson, H. Garcia-Molina, and D. Skeen. Consistency in a partitioned network: a survey. *ACM Computing Surveys (CSUR)*, 17(3):341–370, 1985.

[17] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels. Dynamo: amazon's highly available key-value store. In *ACM SIGOPS Operating Systems Review*, volume 41, pages 205–220. ACM, 2007.

[18] M. Fowler. Nosqldefinition. `http://martinfowler.com/bliki/NosqlDefinition.html`, 2012. [Online; accessed: 2015-02-05].

[19] T. Haerder and A. Reuter. Principles of transaction-oriented database recovery. *ACM Computing Surveys (CSUR)*, 15(4):287–317, 1983.

[20] T. Harter, D. Borthakur, S. Dong, A. S. Aiyer, L. Tang, A. C. Arpaci-Dusseau, and R. H. Arpaci-Dusseau. Analysis of hdfs under hbase: a facebook messages case study.

[21] P. Kieun. All about Two-Phase Locking and a little bit MVCC. `http://www.cubrid.org/blog/cubrid-life/all-about-two-phase-locking-and-a-little-bit-mvcc/`, 2012. [Online; accessed: 2015-03-18].

[22] S. S. Kulkarni, M. Demirbas, D. Madeppa, B. Avva, and M. Leone. Logical physical clocks and consistent snapshots in globally distributed databases.

[23] H.-T. Kung and J. T. Robinson. On optimistic methods for concurrency control. *ACM Transactions on Database Systems (TODS)*, 6(2):213–226, 1981.

[24] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, 1978.

[25] L. Lamport. The part-time parliament. *ACM Transactions on Computer Systems (TOCS)*, 16(2):133–169, 1998.

[26] K. Marzullo and S. Owicki. Maintaining the time in a distributed system. In *Proceedings of the second annual ACM symposium on Principles of distributed computing*, pages 295–305. ACM, 1983.

[27] F. Mattern. Virtual time and global states of distributed systems. *Parallel and Distributed Algorithms*, 1(23):215–226, 1989.

[28] D. L. Mills. A brief history of NTP time: Memoirs of an internet timekeeper. *ACM SIGCOMM Computer Communication Review*, 33(2):9–21, 2003.

[29] D. Ongaro and J. Ousterhout. In search of an understandable consensus algorithm. In *Proc. USENIX Annual Technical Conference*, pages 305–320, 2014.

[30] K. Roebuck. *Extreme Transaction Processing: High-impact Emerging Technology-What You Need to Know: Definitions, Adoptions, Impact, Benefits, Maturity, Vendors*. Emereo Publishing, 2012.

[31] J. Shute, R. Vingralek, B. Samwel, B. Handy, C. Whipkey, E. Rollins, M. Oancea, K. Littlefield, D. Menestrina, S. Ellner, et al. F1: A distributed SQL database that scales. *Proceedings of the VLDB Endowment*, 6(11):1068–1079, 2013.

[32] D. Skeen. Nonblocking commit protocols. In *Proceedings of the 1981 ACM SIGMOD international conference on Management of data*, pages 133–142. ACM, 1981.

Security-Enhanced Linux policy analysis techniques

Filippo Bonazzi
Student number: 472052
filippo.bonazzi@aalto.fi

Abstract

SELinux is an implementation of Mandatory Access Control on Linux, where the Access Control is enforced according to a centralized security policy.

The policy definition language is very flexible and powerful, allowing for extremely fine-grained policy control over the target system; as a result, SELinux policies are usually complex, and their analysis requires expert knowledge and advanced methods and tools.

This paper gives an overview of the main existing tools and formal methods dedicated to the analysis of the SELinux policy, with a focus on deficiencies and areas of possible future improvement.

KEYWORDS: SELinux, policy, analysis, tools, formal methods, information flow

1 Introduction

Security Enhanced Linux (SELinux)[1] is an implementation of Mandatory Access Control (MAC) on Linux. SELinux allows fine-grained control over system resources, defining the permissions in a centralized, administratively-set policy. This centralization makes the policy more practical to analyse compared to traditional UNIX file system permissions, since the analysis can be focused in a single place without traversing the whole file system; however, the finer-grained control results in a considerable policy size, often rendering its analysis daunting.

SELinux policy analysis is usually performed by different parties, with different goals. The **SELinux project** and **Linux distributions** such as RedHat design policies for real-world systems: their primary interest are **practical tools** to generate, parse, modify and verify policies. This has acquired a particular weight in the past few years, after the introduction of SELinux on the Android mobile platform. The problem of policy analysis has also received interest in **academia**: SELinux policy analysis constitutes a fairly high-profile use case for formal verification, where the correctness of a particular policy can be proved with respect to given formal security requirements.

These two groups are not necessarily disjoint, and these efforts exist side-by-side. This paper gives a survey of the main tools and approaches devised in the last 15 years. Section 2 gives a brief introduction on SELinux and the SELinux policy. Section 3 presents practical tools designed to parse, analyse and verify policies; Section 4 presents formal verification tools designed to evaluate an SELinux policy with regard to some given constraints; Section 5 presents some tools devoted to security policy information visualisation. Section 6 presents some more tools related to the topic. Finally, Section 7 identifies some promising areas for future development and provides some concluding remarks.

[1] http://selinuxproject.org/

2 SELinux

SELinux is a Linux Security Module (LSM) [1] providing support for access control policies, usually configured to provide MAC. It was developed in 2001 by the U.S. National Security Agency (NSA), building on previous research projects such as the Flux Advanced Security Kernel (FLASK) architecture and the Fluke research operating system[18]. The main contribution of this architecture was the separation of policy from enforcement: the security logic would be integrally defined in the policy, and a security server would be designed to apply the policy when making security decisions. This conceptual separation has been carried over to SELinux: therefore, it is only necessary to analyse the SELinux policy to obtain a complete description of an SELinux system.

2.1 Security Models

SELinux supports two main security models to regulate access to system resources: **Type Enforcement** (TE) [8][5] and **Role-Based Access Control** (RBAC) [10][19].

It also supports **Multi-Level Security** (MLS) [7][6], which, however, is seldom used, and is therefore left outside the scope of this survey.

All access control can be generalized in terms of subjects, objects, and actions performed by a set of subjects on a set of objects. An access control system controls which actions are allowed by a particular subject on a particular object based on a security policy. The policy describes the manner in which subjects may interact with objects.

In the following sections, this terminology will be used as reference.

2.1.1 Type Enforcement

In SELinux, system resources are categorized in *classes*, which represent the equivalent operating system entities: example classes are *file*, *dir*, *socket*, *process* and more. To control the actions that can be performed on the system resources, *permissions* are mapped to each operating system

ALLOW	**[domain]**	**[type]**	**: [class] {[allowed permissions]}**
allow	my_process	my_file_t	: file {ioctl read getattr lock open}
allow	my_process	my_file_t	: dir {ioctl read getattr search open}
allow	mydomain	my_process	: process {getpgid setpgid dyntransition}

Table 1: Type Enforcement allow rules

primitive: example permissions are *read*. *write*, *open*, *getattr* and more. A given system resource will have a finite set of actions that can be performed on it: the associated *class* will therefore have a finite set of *permissions* that can be granted over it.

System resources are the *objects* in the SELinux security model. In order to perform access control, SELinux has to address *objects* with some stricter granularity than by *class*: to this end, *objects* are assigned a *type* that semantically defines their purpose.

Conversely, processes are the *subjects* in the SELinux security model: in order to perform access control, SELinux assigns them a *domain* that semantically defines their purpose.

Once *subjects* and *objects* are labelled, performing access control is simply a matter of writing *rules* controlling which *domains* are allowed to access which *types*. In detail, an *allow rule* allows a *subject* in a certain *domain* a certain set of *permissions* on an *object* of a certain *class* labelled with a certain *type*.

The format of Type Enforcement rules can be seen in Table 1 along with some sample rules. In the first rule, a process in the `my_process` *domain* is allowed to perform `ioctl`, `read`, `getattr`, `lock` and `open` on a *file* labelled with the `my_file_t` *type*. In the second, another set of permissions (`ioctl`, `read`, `getattr`, `search`, and `open`) is granted on a *dir* labelled with the same `my_file_t` *type*. The third rule allows a process in the `mydomain` *domain* to perform `getpgid`, `setpgid` and `dyntransition` on a *process* labelled with the `my_process` *type*.

In accordance to the principle of safe defaults, SELinux follows a whitelisting approach where every action is denied by default, and explicit rules in the policy grant access to a set of permissions on an object.

2.1.2 Role-Based Access Control

Role-Based Access Control (RBAC) is an abstraction over the concept of "user", where sets of permissions are assigned to administratively defined *roles*; users are subsequently assigned one or more roles, as deemed necessary by the functions they need to perform. This decoupling of users from permissions simplifies user creation and management, and allows for much more concise policy definition and analysis. RBAC is widely used access control model in large corporations[17], since it naturally mimics an organization's behaviour where employees change roles and are dynamically assigned to different tasks over time.

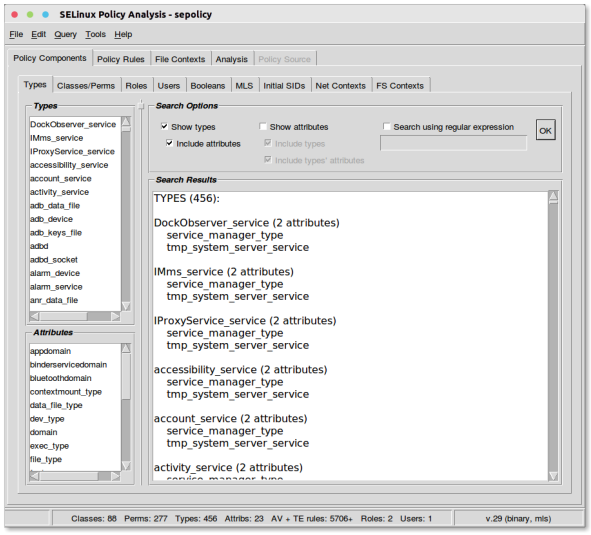

Figure 1: Apol

2.2 Security policy

SELinux security decisions are made according to a security **policy**, which describes the complete state of the security system. The separation of the policy from the Security Server enforcing it is one of the main tenets of the FLASK architecture: the effort of the SELinux research communities can therefore focus mainly on the definition and analysis of the policy, taking for granted the actual implementation of the security architecture and enforcing of the policy.

A SELinux policy is written as a set of plaintext files specifying the entities used in Type Enforcement and Role-Based Access Control: *types*, *domains*, *classes*, *permissions*, *rules*, *users*, *roles*, etc. These plaintext files are compiled into a single binary representation called a **binary policy**: this is done to facilitate efficient policy handling by the kernel when taking security decisions at runtime. SELinux policies are usually distributed in this binary format.

2.3 SEAndroid

SELinux has been ported in the last years to the Android platform: work started in 2011 with a paper by Smalley at NSA [20], and culminated in 2013 [21], when Android 4.4 "KitKat" was released containing a fully functional SELinux port called SEAndroid. However, in Android 4.4 SELinux was configured in permissive mode - actions would be logged but not blocked - and it remained so until Android 5.0 "Lollipop" in late 2014, which featured SELinux in enforcing mode.

While SEAndroid features some additions and differences with respect to SELinux, its basic mechanism and policy definition language are the same. This allows tools and methods for policy analysis to be used across the two domains, the only *caveat* being the different expert knowledge required for each of the two.

The main difference between the two use cases is that SELinux is primarily used on servers, where applications are usually well-known (Apache, Postfix, ...) and individually covered by the policy. In SEAndroid, there is a rather strong distinction between system services and user apps: the former are part of the system, and therefore individually targeted by the policy, while the latter are usually grouped in a common, rather restrictive domain.

SELinux usage has traditionally been somewhat limited, and only specific Linux distributions (RedHat, Fedora) provide an SELinux-enabled system by default. Android "Lollipop" devices are already 5.4% of the total Android devices as of April 2015 [2] : this number is rather significant and only destined to increase, leading to a very widespread usage of SELinux on the Android platform.

3 Practical policy analysis: SETools

Practical policy analysis tools are important to security professionals designing policies for real systems.

The SETools library[3] is a set of tools and libraries for SELinux policy analysis, developed by Tresys Technology. It is the *de facto* standard for handling SELinux policies in text and binary format; the tools it offers are shown in Table 2.

For the sake of brevity, we will give a short overview of the most relevant tools: `apol` and `sediff`.

`apol`	graphical policy analysis tool
`seaudit`	graphical audit log analysis tool
`sediffx`	graphical tool to compare two policies
`seinfo`	tool to get components from a policy
`sesearch`	tool to get rules from a policy
`findcon`	tool to find files matching a context
`replcon`	tool to replace a file's context with another
`indexcon`	tool to index the contexts in a policy
`sechecker`	tool to perform modular checks on a policy
`sediff`	tool to semantically compare two policies

Table 2: SETools tools

3.1 Apol

`Apol` is a graphical tool to extract information from a SELinux policy in several formats. It can interactively query for policy elements, perform information flow and other high-level analyses.

The program is shown in figure 1 performing a Type query.

[2] `https://developer.android.com/about/dashboards/index.html`

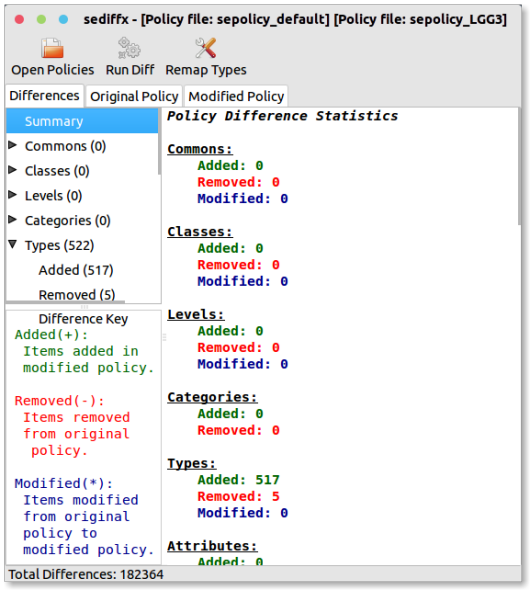

Figure 2: Sediffx

3.2 Sediff/sediffx

`Sediff` is a command-line tool that can find semantic differences between two policies, such as added, removed or modified types, users, roles, and more importantly, rules. This is very useful when analysing policies expanded from a base version, as is the case e.g. in SEAndroid policy development. `Sediffx` is a graphical application that performs the same functions in a graphical environment.

The `sediffx` program is shown in figure 2 reporting some statistics about the semantic difference between two SEAndroid policies.

4 Formal policy analysis

Formal policy analysis tools are often developed as proof of concept of a wider analysis methodology, and do not usually result in actual publicly available software.

4.1 Gokyo

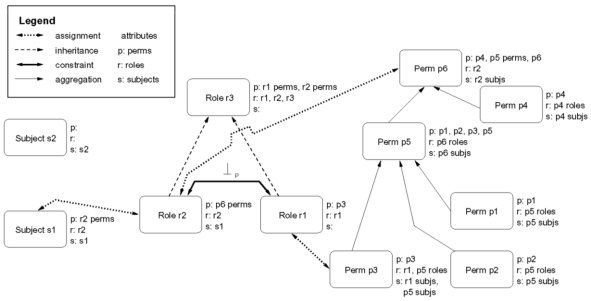

Figure 3: Gokyo example access control model [15]

Gokyo[15] is a policy analysis tool designed to identify and resolve conflicting policy specifications. It has been developed by IBM Research.

The tool takes a higher-level approach to policy design by imposing some constraints on the policy writer to respect the intended security behaviour of the policy. This ensures that the policy specification is correct, and in turn its implementation. The tool helps to prevent problems such as integrity conflicts, contrasting rules, and developer oversights.

Gokyo models a policy as a graph, where Classes, Permissions, Roles and Domains are represented by graph nodes; policy constraints are represented in the form of graph properties. Policy analysis consists of asserting if the corresponding properties hold for the graph generated from the policy under analysis.

An example of an access control specification using this model is shown in Figure 3.

4.2 HRU Formal Verification

The HRU security model is a computer security model which deals with the integrity of access rights in an operating system[13]. It is named after its authors, Harrison, Ruzzo and Ullman, and is primarily used for model safety analysis, with a number of tools and techniques devised for this purpose.

A 2011 paper by Amthor, Kuhnhauser and Polck[4] proposes an approach to SELinux access control policy analysis taking advantage of these techniques. In order to use HRU model safety analysis techniques to analyse SELinux policies, these have to be mapped to an equivalent HRU model first; analysis is then performed on the equivalent model, and the results are used to correct the original policy.

An HRU model is a state machine, where each state is a snapshot of a system's Access Control Matrix (ACM). State transitions happen when applications modify the subject set, object set, or permissions of the ACM. This way, security properties such as right proliferation can be analysed by employing a relatively traditional technique such as state reachability analysis: if a state is reachable starting from a given state with some set of rights, it is termed *not safe* with respect to that given state. Conversely, if a state can never be reached from a given state, it is termed *safe* with respect to that state's set of rights.

The authors propose a procedure to perform the isomorphic mapping of an SELinux policy onto a HRU model, noting that such a procedure has to guarantee *equivalence* - the two systems must behave in the same way after the mapping - and *reversibility* - all the information about the SELinux policy must be preserved in the HRU model. The first major step of this procedure is rewriting the SELinux policy elements as a single composed matrix having the semantics of an ACM. Then the actual HRU model is built on top of this matrix, adding an authorization scheme.

Finally, the authors developed a tool (`sepol2hru`) to automatically perform this model transformation. The tool takes input data from the policy source files containing all the rules, classes and permissions, and from lists which encode an initial system configuration; it outputs the HRU model description in an single XML-based file.

4.3 Information flow

Operating systems security is usually built to guarantee a degree of both data integrity and confidentiality, amongst other requirements. In systems security theory, these two goals can be summed up as *information flow* control: a secure system must be able to control, and possibly restrict, the flow of information between entities.

To preserve integrity, the flow from untrusted sources to trusted destinations must be controlled by a trusted agent able to verify and sanitize the incoming information based on the specific untrusted source. To preserve confidentiality, the flow from trusted sources to untrusted destinations must be controlled by a trusted agent able to sanitize and tailor the outgoing information to the specific untrusted destination.

A 2005 paper by MITRE Corporation researchers Guttman, Herzog, Ramsdell and Skorupka [12] proposes an analysis method for SELinux policies based on information flow analysis. SELinux domain transitions are first mapped into an equivalent information flow diagram, then the the resulting model is checked to determine whether a given integrity or confidentiality goal is enforced by a particular SELinux configuration.

The researchers developed automatic tools to convert the SELinux labeled transition system and a set of diagrams into input for the model checking software NuSMV[3]. The software produces output showing whether the security goal is met by the policy files, and if not, what is the information flow that breaks the security model.

Sample output of the NuSMV model checker is shown in Figure 4.

```
-- specification
!(t = user_t
& E[t != httpd_admin_t U t = httpd_sys_script_t
& EF (k = TRUE & t = httpd_sys_script_t)])
is false
-- as demonstrated by the following execution sequence
-> State 1.1 <-
t = user_t
r = system_r
u = system_u
c = netif_c
p = rawip_send_p
k = 1
-> State 1.2 <-
t = netif_ipsec2_t
r = object_r
u = jdoe_u
p = udp_recv_p
-> State 1.3 <-
t = dpkg_t
r = system_r
u = system_u
c = fifo_file_c
p = append_p
-> State 1.4 <-
t = httpd_sys_script_t
r = object_r
u = jdoe_u
c = netif_c
p = accept_p
```

Figure 4: NuSMV sample output [12]

[3] http://nusmv.fbk.eu/

5 Policy visualisation

SELinux policy analysis is made more complex by the fact that it requires accurate, in depth understanding of the relationships between apparently distant rules, potentially hidden by the imposing size of the policy. Even an expert eye would benefit from an easier way to analyse and browse a SELinux policy: time and effort saved on exploring the policy could be better spent analysing the most sensitive areas. The research community has naturally taken this direction by using extensive procedure automation, as presented in the previous sections. Another fundamental technique, destined to become more and more relevant as SELinux-enabled systems become more widespread, is data visualisation. Data visualisation allows key policy information to be understood visually, without reading through high amounts of text, and allows higher-level abstractions to be used in policy analysis.

5.1 SPTrack

A 2012 paper by Clemente, Kaba, Rouzaud-Cornabas, Alexandre, and Aujay [9] presents a method and tool (`SPTrack`) to allow the manipulation of real SELinux security policies and to visualize potential security violations.

The researchers represent a SELinux policy as a graph, where nodes are SELinux objects (including subjects) and edges are the interactions between the objects; SELinux classes and permissions are associated with each edge of the graph.

Given the somewhat unwieldy size of such a visualisation, they also present a reduced graph displaying only information flows. In this second graph the researchers resort to a notion of *criticality* to classify flows; they informally define criticality as "*the potential erasability provided by the interaction (syscall) for the subject to alter the integrity or the confidentiality of the object*" [9]. Following this definition, flows are assigned the colours green, blue, yellow and red by increasing criticality.

These graphs are realized with the `SPTrack` tool developed for the purpose. The two kind of graphs - full policy and information flow - are shown in Figure 5. The information flow graph on the right has been reduced to show only medium and high criticality flows.

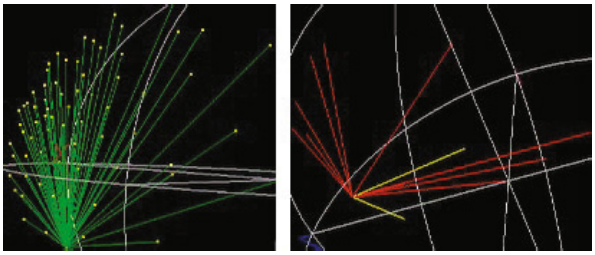

Figure 5: Full policy and information flow graph [9]

5.2 SEGrapher

A 2011 paper by Marouf and Shehab [16] proposes a clustering technique and visualisation tool (`SEGrapher`) to analyse SELinux policies.

The authors propose a clustering algorithm that groups SELinux types based on their allowed permissions. Clusters provide a simple abstraction that allows humans to visually discover interesting and inherent relations between types, without getting lost in the complexity of the policy definition.

The `SEGrapher` tool implements the proposed clustering algorithm and provides a way to visualize the clusters and their interactions. The tool is based on the Java JDK 1.6, and uses the SETools libraries (Section 3) for parsing SELinux policies. A sample cluster graph generated by `SEGrapher` is shown in Figure 6.

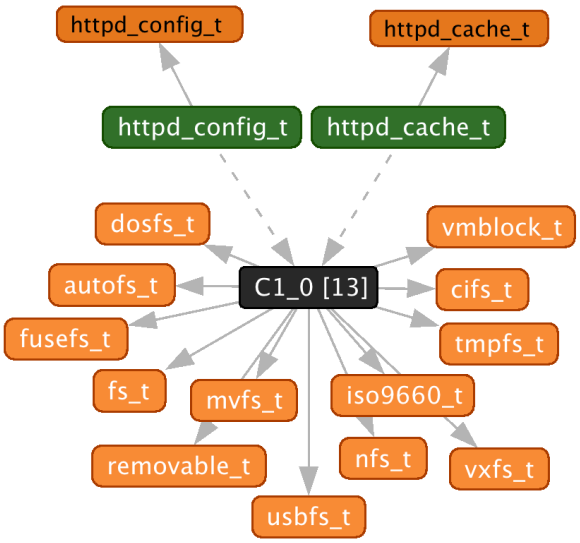

Figure 6: SEGrapher cluster graph [16]

6 Related work

Many more tools have been written to work with SELinux policies: we give here a brief overview of some tools that are somewhat relevant to the topic, but do not deserve a more in-depth presentation.

6.1 Polgen

Polgen is a tool for semi-automated SELinux policy generation, developed by the MITRE corporation [22]. It is aimed at system administrators - who may be tasked with integrating unfamiliar applications into existing SELinux policies - and at application authors - who may want to build policies for their applications without being overly experienced in SELinux.

Polgen appears to be no longer in active development, and most authoritative sources are no longer available on the Internet.

6.2 audit2allow

Audit2allow is a tool for automated SELinux policy generation, developed in 2006 by the NSA as part of the SELinux project userspace tools [2]. Using audit2allow policy writers can convert SELinux audit messages into rules, allowing quick and easy policy generation for unfamiliar applications. The generated policy, however, is not necessarily correct, complete or secure, since it is not designed to be so: the rules depends on code paths taken by the execution flow, and there is no way to distinguish intended and possibly malicious application behaviour.

6.3 Shrimp

There has been some research in applying Domain Specific Languages (DSL) [11] to SELinux policy development and verification.

A 2009 paper by Hurd *et al.* [14] proposes two Domain Specific Languages, Lobster and Symbion, to specify SELinux policies at a high abstraction level, and then render them more specific without loss of the original meaning through a process termed *refinement*. In the Lobster language, a policy is represented as a program: when interpreted, the program generates the equivalent information flow graph, from which the SELinux policy is in turn generated. The Symbion Language is an assertion language for information flows in Lobster policies. Symbion flow assertions are attached to Lobster policies: when the Lobster policy is interpreted into the equivalent information flow graph, the resulting flows are checked against the assertions, ensuring the correct behaviour of the model.

The researchers then propose a tool (`shrimp`) to analyze and find errors in the SELinux Reference Policy, much akin to the Lint tool for the C programming language. `Shrimp` outputs HTML policy documentation and analysis results.

7 Conclusions

The SELinux policy analysis field deals with a complex and sensitive subject, which is becoming more and more relevant as SELinux adoption increases across desktop and mobile platforms. This increasing importance is bound to cause larger involvement by corporations - such as security companies and Android OEMs - in policy development: this will indubitably lead to the development of better and more advanced tools to deal with SELinux policies, from the original design to quality assurance and security analysis.

Some of these improvements are already taking place today: the SETools library presented in Section 3 is currently being rewritten in Python as SEToolsv4, and also being expanded to better deal with the SEAndroid use case.

The greatest developments will come in the area of automated policy security analysis, where it is evident the lack of a tool to perform heuristic policy analysis and reporting. Such a tool, working on an expert-provided knowledge base of anti-patterns and common mistakes, would prove invaluable in SELinux policy quality assurance; it would also be an interesting first step towards checking whether a policy is *minimal* or not with regard to its intended behaviour.

Future developments of such a tool could also include higher-level policy verification, as proposed by Gokyo in Section 4.1 and Shrimp in Section 6.3, to guarantee that the actual policy is *correct* and *complete* in line with its original specification.

Significant developments are also to be expected in the policy visualisation area, most likely as built-in additions to existing policy analysis tools such as those in the SETools library.

References

[1] Linux Security Modules Kernel Documentation. `https://www.kernel.org/doc/Documentation/security/LSM.txt`. Accessed: 2015-04-06.

[2] SELinux Userspace. `https://github.com/SELinuxProject/selinux/wiki`. Accessed: 2015-04-06.

[3] SETools Policy Analysis Tools for SELinux. `https://github.com/TresysTechnology/setools3/wiki`. Accessed: 2015-02-06.

[4] P. Amthor, W. Kuhnhauser, and A. Polck. Model-based safety analysis of selinux security policies. In *Network and System Security (NSS), 2011 5th International Conference on*, pages 208–215. IEEE, 2011.

[5] L. Badger, D. F. Sterne, D. L. Sherman, K. M. Walker, and S. A. Haghighat. Practical domain and type enforcement for UNIX. In *Security and Privacy, 1995. Proceedings., 1995 IEEE Symposium on*, pages 66–77. IEEE, 1995.

[6] D. E. Bell and L. J. La Padula. Secure computer system: Unified exposition and multics interpretation. Technical report, DTIC Document, 1976.

[7] D. E. Bell and L. J. LaPadula. Secure computer systems: Mathematical foundations. Technical report, DTIC Document, 1973.

[8] W. E. Boebert and R. Y. Kain. A practical alternative to hierarchical integrity policies. *NIST SPECIAL PUBLICATION SP*, pages A–10, 1989.

[9] P. Clemente, B. Kaba, J. Rouzaud-Cornabas, M. Alexandre, and G. Aujay. Sptrack: Visual analysis of information flows within selinux policies and attack logs. In *Active Media Technology*, pages 596–605. Springer, 2012.

[10] D. Ferraiolo, D. R. Kuhn, and R. Chandramouli. *Role-based access control*. Artech House, 2003.

[11] M. Fowler. *Domain-specific languages*. Pearson Education, 2010.

[12] J. D. Guttman, A. L. Herzog, J. D. Ramsdell, and C. W. Skorupka. Verifying information flow goals in security-enhanced linux. *Journal of Computer Security*, 13(1):115–134, 2005.

[13] M. A. Harrison, W. L. Ruzzo, and J. D. Ullman. Protection in operating systems. *Commun. ACM*, 19(8):461–471, Aug. 1976.

[14] J. Hurd, M. Carlsson, S. Finne, B. Letner, J. Stanley, and P. White. Policy DSL: High-level Specifications of Information Flows for Security Policies. 2009.

[15] T. Jaeger, R. Sailer, and X. Zhang. Analyzing integrity protection in the selinux example policy. In *Proceedings of the 12th conference on USENIX Security Symposium-Volume 12*, pages 5–5. USENIX Association, 2003.

[16] S. Marouf and M. Shehab. SEGrapher: Visualization-based SELinux policy analysis. In *Configuration Analytics and Automation (SAFECONFIG), 2011 4th Symposium on*, pages 1–8. IEEE, 2011.

[17] A. C. O'Connor and R. J. Loomis. 2010 Economic Analysis of Role-Based Access Control. *NIST, Gaithersburg, MD*, 20899, 2010.

[18] N. Peter Loscocco. Integrating flexible support for security policies into the linux operating system. In *Proceedings of the FREENIX Track:... USENIX Annual Technical Conference*, page 29. The Association, 2001.

[19] R. S. Sandhu. Role-based access control. *Advances in computers*, 46:237–286, 1998.

[20] S. Smalley. The case for SE Android. *Linux Security Summit*, 2011.

[21] S. Smalley and R. Craig. Security Enhanced (SE) Android: Bringing Flexible MAC to Android. In *NDSS*, volume 310, pages 20–38, 2013.

[22] B. T. Sniffen, D. R. Harris, and J. D. Ramsdell. Guided policy generation for application authors. In *SELinux Symposium*, 2006.

Bacteria Nanonetworks

Erik Berdonces Bonelo
Aalto University School of Science
erik.berdoncesbonelo@aalto.fi

Abstract

As nano technologies are being developed, nano-scale devices can execute more tasks and a new field appears: how to coordinate and interconnect swarms of nanodevices. The progress in this field will allow new solutions in nanomaterial creation and drug delivery.

This article describes new fundamental methods to create nanonetworks and the major related challenges. While there are many different approaches, we specifically focus on bacteria conjugation and plasmid transmission as foundations of bacteria communications. This paper provides an exhaustive description of the different steps for transmitting through bacteria information and allows to see the evolution and future of the nanonetoworks technology as well as the uses of it.

KEYWORDS: bacteria, nanonetworks, DNA, nanodevices

1 Introduction

In 1959, the Nobel laureate physicist Richard Feynman gave his famous speech titled "There's Plenty of Room at the Bottom", in which he detailed the paradigm of working with individual atoms and molecules. Nowadays, nanotechnology is an emerging field that is in starting to grow. In the last decade it has enabled the creation of nanoscale devices which allow new procedures in the healthcare, industrial and military fields. These nanomachines can do many tasks, such as encapsulating drugs and delivering them to specific areas, as well as building intelligent materials.

However, recent research has pointed out the need to develop systems to interconnect these swarms of nanodevices. The reason comes from the fact that traditional techniques are not applicable at nano scale [2]. Different techniques are being developed in order to solve this new paradigm of nanoscale exchange of information, however this article especially focuses on networks of nanodevices engineered by bacteria. The main idea is to use the communication between bacteria to transmit the desired data to the desired node.

This article, first, explains the motivations to investigate bacteria nanonetworks and nanotechnology itself. Second, describes the different possible architectures in bacteria nanonetworks. Next we will discuss the performance and benefits of each design. Finally, we will describe the main issues and the obstacles to overcome in order to make this technology feasible.

2 Motivation

Recent developments in nanotechnology have allowed the creation of machines in a nanoscale. The nanoscopic scale (or nanoscale) usually refers to structures with a length scale applicable to nanotechnology, usually cited as 1âĂŞ100 nanometers, a billion times smaller than of a meter. The nanoscopic scale is (roughly speaking) a lower bound to the mesoscopic scale for most solids.There is a huge variety of areas where nanomachines can be useful. Use-case examples are new vaccines, biological and chemical sensors, development of new manufacturing processes and drug delivery systems. The creation of these small machines has been possible as a result of the development of production techniques of new materials, such as graphene, whose characteristics allow us to operate in a nanometric scale both easily and efficiently.

The structure of these nanomachines is based on naturally occurring bacteria and can interact with them. Moreover, they can include DNA parts which interact and can relay and retrieve DNA information. Also, they can include components that facilitate mobility or even small graphene antennas which work over the terahertz band [3].

However, the main problem is the communication and coordination between the different individual nanomachines. Even if it is possible to use graphene antennas over a terahertz bandwidth, the drawback is that nanomachines do not possess a convenient energy source, and molecules and bacteria could still interfere in the signal transmission, as they are not negligible if compared to the wavelength.

Because of these limitations, scientists have looked at other approaches. The main one they are developing consists in using the systems that nature has already created: molecules. The inspiration to use molecular communications is drawn from the communications existing in nature [9]. The approach is to code the message in the change in the value of the concentration of certain molecules (signals) in the medium. This system is used in short range communications between bacteria.

Unfortunately molecular based communications have huge drawbacks, such as very low capacity or the need for a complex infrastructure [8]. Also, in long range communications between bacteria, this technique is not useful. In order to overcome this issues, scientists have started to use bacteria to carry messages encoded in their DNA. These bacteria move from the transmitter to the receiver nodes and could be routed as in traditional data networks [6].

3 Background

3.1 Bacteria motility and propagation through chemotaxis

Motility is the term used to define mobility in bacteria and other unicellular and simple multicellular organisms. They can move by slugging or using cilia and flagella. Cilia are short appendages extending over a surface of the cell. Flagella are long appendages which work by rotating as a propeller to thrust the cell. Different bacteria have different motility speeds depending on the method used to propel themselves. They can reach speeds ranging from $2\,\mu\mathrm{m/s}$ (*Beggiatoa*) to $200\,\mu\mathrm{m/s}$ (*Vibrio Cholerae*) [6].

Bacteria choose their motility direction through a sensing process called *chemotaxis*. They sense the concentration gradient of some specific beneficial chemicals (called chemoattractant) in their surroundings and alter their motility path according to it. For instance, *E. coli* adapts a random walk when there is no chemoattractant in effect. However, within a chemoattractant field, it will move in a biased random walk model towards the attractant [14]. Depending on the chemoattractant, the random walk bias is different.

Bacteria can also move as a swarm. When moving as such, bacteria group and help each other to advance and propagate. Depending on the bacteria species, the propagation of a swarm can affect the swarm itself, limiting the direction of propagation.

3.2 Plasmids

Plasmids are circular DNA strings which bacteria hold inside them. They must not be confused with the chromosomal DNA as they are contained in a different molecule. This plasmid can replicate independently of the bacterium nucleus (self-replication) and eventually transfer the encoded information to other bacteria (self-transfer). The transfer of the plasmid to other bacteria takes place through hairlike appendages called pili. These appendages connect both bacteria and transfer the genetic material in the plasmid. This process, called bacteria conjugation, allows the interchange of information between bacteria.

The DNA of a plasmid consists of a double-stranded DNA molecule which comprises two polymers nucleotides. Each nucleotide contains four possible bases: adenine (A), cytosine (C), guanine (G) or thymine (T) [8]. Each of them determines the base of the other nucleotide in the pair, thus the combinations of the DNA are either AT or CG. Each couple of nucleotides forms a *base pair* (bp), thereby encoding 2 bits of information. Plasmids can be up to 1.6 mbp (mega base pairs) long [8].

Plasmids can be transfered to species different from that of the original information carrier. To do this, both bacteria need to be in contact. Also, through the transmission of plasmids, special characteristics of a bacteria can be shared with other bacteria, such as vulnerability to a certain antibiotic or fluorescence. This paper focuses on the transmission of the coded information in the DNA.It is important to note that the success rate of the transmission depends on the specific species combination. Usually, the conjugation between bacteria of the same species is the combination that has the best probability to be performed.

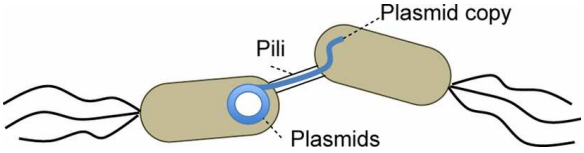

Figure 1: Illustration of the conjugation process involving two bacteria. A copy of the plasmid is transferred through the connection provided by the pili.

3.3 Nanomachines as nodes

Nanomachines are functional devices made of nanoscale components and can perform certain tasks. An interconnected swarm or group of them can form a nanonetwork. Nanomachines can be created by using man-made components or reusing biological entities found in nature. There are three approaches to creating nanodevices [4]:

- *Top-down*: nanomachines are developed through downscaling current microelectric technology to a nanoscale size.
- *Bottom-up*: realized through the assembly of molecular and synthesized nanomaterials.
- *Bio-hybrid*: developed man-made nanomachines using the previous methods are combined with existing biological components from bacteria.

The technology for man-made nanomachines has been evolving due to the development of classical lithography and the study and production of nanomaterials. However, at the current state of the art, it is more feasible to use or copy components already made by nature. One example is using Adenosine TriPhosphate or ATP batteries, inspired by the behavior of *mitochondria*.

In bacteria nanonetworks, nanomachines could be used as nodes that transmit and receive from and to where the information is routed to. A DPU (Data Processing Unit) component for nanomachines is being developed, which can encode and decode strings of DNA, such as a DNA-based Turing Machine [7].

3.4 Actual challenges

Current challenges in the field are related to the different behavior that bacteria exhibit in different environments. Example of problems include the selfish behavior when different breeds of bacteria encounter in a scenario having a low level of nutrients [11]2. Some bacteria will try not only to kill off other species but also members of the same species. The density of bacteria population can also trigger this behavior switching and thus stop cooperation. Part of this behavior can be cut down using antibiotics that could be dispensed by nanomachines located in the environment. However, this raises another problem: bacteria can increase their resistance

to these antibiotics and can also obtain immunity to some antibiotics as a result of bacteria conjugation.

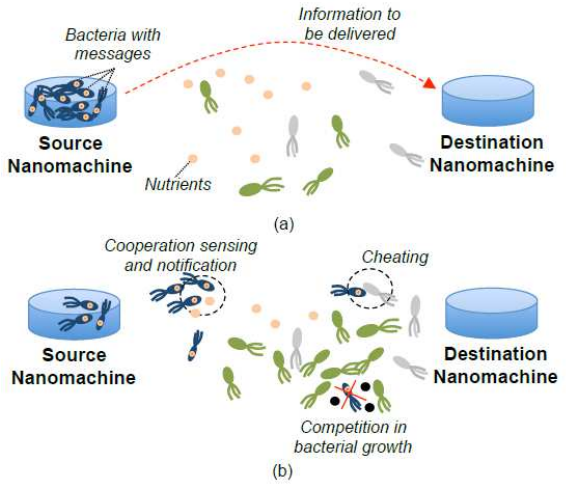

Figure 2: a) illustrates bacteria behaviour in coopertive mode. In contrast b) illustrates a case with a non cooperative behaviour, where bacteria cheat and compete.

Another challenge related to nanomachines is the assembly and grouping of the different components in them. Even if the different pieces forming the nanomachines can be created with our actual technology, we are still limited in moving them to the proper position so that they can perform their function inside organ tissues for instance.

Solution to these challenges already exist, such as using antibiotics to kill undesired bacteria. However, better and more efficient approach should be found.

4 Communication Techniques

New investigations in bacteria communication technologies have developed different solutions. This section covers the most relevant ones, describing the basic operations as well as the supporting and opposing arguments for each of them.

4.1 Molecular signaling: diffusion

Molecular signaling is a technique used by cells to communicate each other through the diffusion of molecules. Cells close to each other deploy specific kinds of molecules to warn or communicate with surrounding cells. Each type of molecule symbolizes a different message. Also, the gradient of concentration of the molecule can indicate different messages or rhe intensity of an event. An example is given by calcium signaling in intracellular communication, where calcium molecules travel through the small gap junction between adjacent cells.

Another example is given by engineered organs in-a-chip. In a small petri dish with interconnected cavities, each cavity can hold different bacteria. The communication between the different cavities or parts of the organ is achieved through the connecting channel. This system allows, among other uses testing of the effects of bacteria in the digestive tract.

Diffusion is frequently used for bacteria communication in laboratories. However, it has some drawbacks in extensive bacteria nanonetwork. First, transmission speed and range is limited. Second, there is no control over the amount information or its routing because it is encoded in diffused molecules. Finally, diffusion relies on a liquid medium in which to diffuse the signaling.

One sample use of diffusion is in nanofulidics, where bacteria are isolated by a membrane composed of a large number of parallel nanocapillaries. Bacteria are not able to move by through the membrane but they still can communicate using diffusion because molecules are small enough to move across the membrane's nanocapillaries.

4.2 Molecular motors communication

One existing solution for intracellular communication employs molecular motors. This approach relies on a grid of tracks made of cytoskeleton. A cytoskeleton is a network of fibers composed of proteins realizing tracks that can be organized in star or mesh topologies. Then, kinesin and myosin motors [13] can be used to carry molecules between the nodes. This system mimics actual wires used to connect computers and allows point-to-point intracellular communication. However, the drawback of this solution is that it cannot be applied to inter-cellular communications.

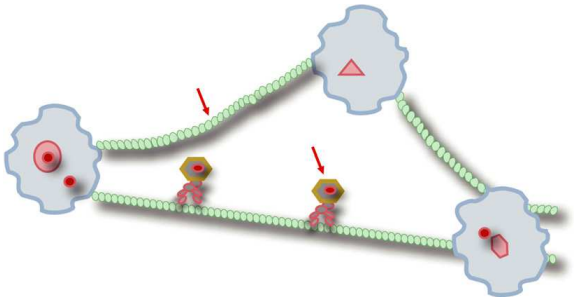

Figure 3: Intracellular comunications using molecular motors to carry molecules between nodes.

4.3 Bacteria conjugation

Bacteria conjugation is based on the transmission of data between bacteria in form of plasmids. As an overview, bacteria transport plasmids that have encoded both information and receiver information. Nodes attract the bacteria diffusing chemoattractants in the medium. An extension of this system can even allow a multi-hop architecture where the information is carried by the same bacteria. This system is detailed next.

5 Conjugation-based nanocommunication

This section details bacteria nanonetworks based on conjugation. The first steps are executed before the bacteria transmission and consist in setting up the address information in the nanomachine nodes. Then we describe the encoding, encapsulation, propagation and relay process using the bacteria and nanomachines described in the previous section.

The use of bacteria conjugation and coding of the information in the DNA allows higher throughput than other methods described in section 4.

5.1 Growing bacteria

The first step in building bacteria nanonetworks is choosing which kind of bacteria will be carrying the information. Each bacterium has different specifications in terms of motility and difficulty in loading and programming their behavior. Another option is using different species at the same time, but this can end in bacterias killing each other.

Bacteria can mutate and deviate of their programmed behavior. For example, *E. coli* has a mutation rate of 10^{-8} per base pair per generation. These kinds of bacteria are rare but still they can be present in a large population.

One of the solutions to avoid bacteria carrying wrong information or that are disobedient is using antibiotics. Bacteria can be programmed to be resistant to a certain type of antibiotic enabling some specific plasmids in the bacteria. Thus antibiotics only affect other non desired bacteria which do not have such resistance, which carry undesired mutations and information. Furthermore, this technique reduces the "noise" of the information, thus increasing the amount of bacteria in the population that carry the same information. One of the locations where to deploy the antibiotics are the nanomachines. The nodes that attract bacteria to be used as carriers kill at the same time the ones that do not fulfill the requirements to relay the information.

Besides its benefits, this procedure also creates a problem: bacteria can develop resistance to a certain antibiotic. If bacteria develop this mutation, the chosen antibiotic can be ineffective. This is one of the reasons to choose a specific breed of bacteria when performing the transmission of data.

5.2 Nanomachine address configuration

Nanomachines use *remote attractants* to attract the bacteria carriers to them, and thus guiding the message towards them. Instead, a node willing to transmit information, releases a *transmission attractant* that will attract the surrounding bacteria to transfer information to.

Nanomachines use an address configuration similar to traditional information networks. They are identified with a two tier address system, consisting in a *physical address* and a unique *network address*.

The network address contains the information of the final receptor of the information. The physical address contains the information of the chemoattractant of the next node. Because a remote attractant can be reused and thus a physical address can be duplicated, the network address removes this ambiguity.

Each node contains a routing table that matches addresses with remote attractants. Every time an information carrier bacteria reaches the node, the physical address is updated to the new remote attractant found in the routing table, thus it points to the next hop in the multi-hop architecture.

5.3 Encoding

A nanomachine node encodes the information into a DNA string which can be transmitted in plasmids. This plasmid can be divided in three basic parts:

- *Transfer region*: region that contains the genes that stimulate the plasmid to self-replicate and transfer.
- *Routing region*: contain the genes which make the bacteria able to carry the message. This genes allow to: deactivate the chemotaxis towards the transmitters and enable it towards the receiver (update the physical address), inhibit bacterial replication (to control the size of the swarm), and enable a programmed death on timeout (that could be translated as the time to live of the bacteria) which prevents delivery of messages with a long delay [8].
- *Message region*: it contains the network address and the body of the message to be transmitted.

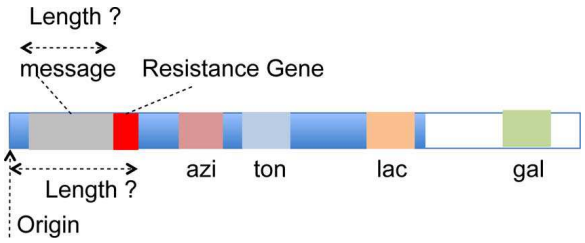

Figure 4: Encoding of a message in a plasmid through the combination of different proteins.

The encoding also poses a new problem which is that some sequences enable some functions inside the bacteria rather than containing a message. This randomly created proteins can interfere with the message delivery and disable some of the expressions in the routing region. So, when encoding the message, these specific sequences have to be avoided. Such sequences may vary between different bacteria.

5.4 Encapsulation and propagation

In this step, the plasmid generated as the node in the previous step is transferred to the bacteria. Empty carriers are attracted to a node by using the transmission attractant. Then, through bacteria conjugation, the plasmid is transferred to each bacteria. Once the conjugation is complete, the active part of the plasmid changes the behavior of the carrier in order to match the one described in Section 5.3.

Propagation starts after the configuration of the bacterium. If bacteria do not reach the receiver during the Time To Live (TTL) recorded in the plasmid, the bacterium kills itself. We assume that a bacterium dies by breaking its membrane and releasing its content to the medium. This poses a new problem, which is that other bacteria can absorb such a content, including the DNA plasmids, in a process called competence. To avoid this problem, non-competent bacteria should not be used when performing transmissions.

If a bacteria die during its propagation, this could be counted as a message loss probability.

5.5 Decapsulation and Relaying

To transfer the information from the bacteria to the node, again, bacteria conjugation is performed. However, a bacterium can try to re-deliver a message, thus the node will kill the bacterium right after the delivery. If several bacteria try to transmit the same message, message duplication may result. The node can exclude duplicated messages in order that it has only one copy of the plasmid. This mechanism is already used by bacteria to avoid accumulating the same plasmid several times.

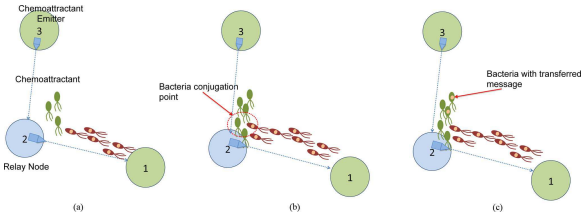

Figure 5: a) bacteria travel to the remote node by following the chemoattractant. b) the node loads new routing information, thus bacteria will follow a new chemoattractant. c) bacteria travel to their new destination

In case several nodes with the same attractant exist in the medium, they can be identified by their physical address thus only the real receiver will accept the message. In some cases, the nodes can be mobile, therefore protocols for opportunistic routing must be used [6]

After decapsulation, the receiver checks the network address. If it matches the local address, it decodes the message. Otherwise, the plasmid is refactorized thus the physical address matches the next hop in the transmission depending on what the node's routing contents [1]. The final node will then decode the message and in turn perform the desired actions with the information.

5.6 Performance and analysis

When using *E. coli*, some of the relevant figures [8] are:

- *Delay*: the estimation is of 10 minutes per relay.
- *Capacity*: it is determined by the replication rate of the bacteria and probability of loss of the carriers.
- *Range*: maximum range of each hop is tens of centimeters.

Overall, the technology is limited as the speed is not very high, but for small networks and nano level communications it can be really useful.

6 Conclusion

Bacteria nanonetworks are a technology still in development and with many issues that have to be solved but that will have a great impact in the future. Besides, scientists are developing architectures and systems that would enable the future commercialization and global use of the technology.

However, hardware is still under development and there are other challenges and paradigms to be solved. There are several uses aside networking and information transmission, such as drug delivery and creation of intelligent nanomaterials even if the transmission delay is extensive and the range is limited.

In summary, bacteria nanonetworks will enable the possibility of using the technology that nature has developed during millennia in a productive way for industry and health.

References

[1] J. Adler. Chemotaxis in bacteria. *Science*, 153(3737):708–716, 1966.

[2] I. F. Akyildiz, F. Brunetti, and C. Blázquez. Nanonetworks: A new communication paradigm. *Computer Networks*, 52(12):2260–2279, 2008.

[3] I. F. Akyildiz and J. M. Jornet. Electromagnetic wireless nanosensor networks. *Nano Communication Networks*, 1(1):3–19, 2010.

[4] I. F. Akyildiz, J. M. Jornet, and M. Pierobon. Nanonetworks: A new frontier in communications. *Communications of the ACM*, 54(11):84–89, 2011.

[5] S. Balasubramaniam et al. Opportunistic routing through conjugation in bacteria communication nanonetwork. *Nano Communication Networks*, 3(1):36–45, 2012.

[6] S. Balasubramaniam and P. Lio. Multi-hop conjugation based bacteria nanonetworks. *NanoBioscience, IEEE Transactions on*, 12(1):47–59, 2013.

[7] Y. Benenson, T. Paz-Elizur, R. Adar, E. Keinan, Z. Livneh, and E. Shapiro. Programmable and autonomous computing machine made of biomolecules. *Nature*, 414(6862):430–434, 2001.

[8] L. C. Cobo and I. F. Akyildiz. Bacteria-based communication in nanonetworks. *Nano Communication Networks*, 1(4):244–256, 2010.

[9] M. Gregori and I. F. Akyildiz. A new nanonetwork architecture using flagellated bacteria and catalytic nanomotors. *Selected Areas in Communications, IEEE Journal on*, 28(4):612–619, 2010.

[10] M. Gregori, I. Llatser, A. Cabellos-Aparicio, and E. Alarcón. Physical channel characterization for medium-range nanonetworks using flagellated bacteria. *Computer Networks*, 55(3):779–791, 2011.

[11] M. Hasan, E. Hossain, S. Balasubramaniam, and Y. Koucheryavy. Social behavior in bacterial nanonetworks: Challenges and opportunities. *arXiv preprint arXiv:1411.4214*, 2014.

[12] G. J. d. A. Soler-Illia, C. Sanchez, B. Lebeau, and J. Patarin. Chemical strategies to design textured materials: from microporous and mesoporous oxides to nanonetworks and hierarchical structures. *Chemical Reviews*, 102(11):4093–4138, 2002.

[13] R. D. Vale and R. A. Milligan. The way things move: looking under the hood of molecular motor proteins. *Science*, 288(5463):88–95, 2000.

[14] Z. Wang, M. Kim, and G. Rosen. Validating models of bacterial chemotaxis by simulating the random motility coefficient. In *BioInformatics and BioEngineering, 2008. BIBE 2008. 8th IEEE International Conference on*, pages 1–5. IEEE, 2008.

Survey on indoor localization methods using radio fingerprint-based techniques

Christian Cardin
Student number: 397179
christian.cardin@aalto.fi

Abstract

The need for precise location-aware applications is increasing along with the technological advancement of the current generation of mobile devices not only for user navigation, but also for accessing services based on location or automatic machine motion. Due to GPS limitations, satellites cannot be used for indoor navigation purposes; for this reason, various indoor positioning solutions have been studied and proposed in recent years. One of the most popular techniques is the localization based on radio fingerprint, which exploits the existing radio infrastructure in the building (such as WLAN, Bluetooth or other mediums), offers sufficient precision and does not require additional costs. This paper surveys the currently available technology for indoor localization based on fingerprint, including methodologies for data collection, algorithms for fingerprint pattern recognition and methods for reducing prediction errors.

KEYWORDS: indoor positioning system, location techniques, fingerprint, WLAN

1 INTRODUCTION

Indoor localization using fingerprinting is part of a larger group of methodologies called Indoor Positioning System (**IPS**), utilized for inferring the position of a device inside an indoor environment, for example a building[8]. Many popular applications available for smart devices take advantage of the user's position to offer location-based services such as navigation (Google Maps), finding the closest point of interest (Foursquare) or knowing friend's activities in the area (newly developed Finnish app NAU). All these services require geo-coordinates from the device's GPS sensor, which are generally not available in indoor areas without a direct line of sight to the satellites. An Indoor positioning system could overcome the problem of the reachability by offering a localization service where the standard GPS is not normally available due to environmental conditions, plus increasing the accuracy down from several meters to few centimetres[10]. In addition, it could provide the basis for a completely new set of services that target the indoor environment offering important benefits to fields like healthcare, accessibility, safety and machine automation. A simple smartphone application could give guidance to blind people for walking inside an hospital; also, it can help to track personnel or equipment inside a large facility and enable robots to move autonomously across rooms and corridors.[13]. Drones and auto-piloted robots are gaining popularity and in the next future they will be used for logistics and transportation.

IPSs can work with different wireless technologies, including IR, RFID, ultra-sound, Bluetooth, WLAN and magnetic field. Every medium offers unique advantages and drawbacks in performing location sensing. The majority of already available solutions utilize the building's WLAN network because it does not require additional costs for a dedicated infrastructure and WiFi is almost everywhere. Another difference that distinguishes indoor localization technologies is the type of algorithm and procedure used to estimate the position: triangulation, fingerprint, proximity and vision analysis.

Triangulation is based on the geometric property of triangles that, knowing the coordinates of three points in the space, makes it possible to infer the position of a fourth point by knowing the distances or the directions from all the other points. In practice, the reference points are known wireless Access Points (AP), and the unknown point (the device) calculates the distance from them by measuring the time of arrival (TOA) of a special ping message[8]. Alternatively the measuring device could use the Angle of Arrival (AOA) which requires only two known points but it is generally less accurate than the previous method, because the signal may bounce on the walls compromising the result[10].

Proximity localization uses RFID tags and it is useful for determining the presence or not of a target inside a limited area, such a room. A certain number sensors are placed in known positions, and when the device enters the range of one sensor the system maps the device's location with the sensor's location.[2]. The vision-based localization uses pictures recorded by the device's camera to recognize known places in a database, or the user can be tracked by known cameras using face recognition techniques[5]

Indoor localization method based on radio fingerprints is usually preferred because it does not need dedicated hardware, it is easy to implement, requires less processing power and has good accuracy. It works in two separate phases: the offline phase and the online phase. The first phase consists of creating a map of the building collecting fingerprints from known locations and saving them in a database: for every reachable AP at the location, collect the ID and Radio Signal Strength Indicator (RSSI), combine them together and save them in the database along with the location information. The Online phase consists of collecting a fingerprint from

an unknown position and checking it against the database. If the fingerprint is registered in the database, the user can be tracked. A drawback is that if the network topology changes the radio map must be built again, thus this method is not suitable for dynamic environment where the APs change often.

This paper introduces the technique of indoor localization using fingerprints collected from the area of interest, compares its characteristics and poses attention on the differences in different implementations in terms of sensing medium and algorithms. Then, this paper shows a typical architecture for an indoor localization service using fingerprinting, including all the intermediate phases and challenges for the initial data collection. Finally, there are considerations of some challenges: error mitigation, security and efficient data collection.

2 CATEGORIES of FINGERPRINT LOCALIZATION TECHNIQUES

This section compares the most common algorithms and technologies used in fingerprinting localization techniques. There are several attributes used for comparing different localization techniques. These are: security, cost, performance, robustness, complexity, user preference, availability and limitations [8].

2.1 Mediums

The fist way to distinguish between fingerprint techniques is to look at the sensing medium used to collect fingerprint values.

A) **WLAN:** The most common sensing medium used in indoor localization is the existing WLAN infrastructure, because it usually covers most part of the building and it does not involve additional costs for installation. Many existing indoor localization services[12], such as RADAR, COMPASS, Ekahu, use the WLAN for its convenience and availability. The localization accuracy varies greatly depending on the geometry of the area and the algorithms used, but it ranges from 1 meter to 5 meters. As explained in section 2.2, some algorithms rely on the time the signal takes to reach the user; if there are walls between the AP and the user, the signal can bounce and the measured time of arrival can be longer than if there is a direct line of sight, resulting in a wrong distance estimation and, as a consequence, an incorrect localization. Generally, the implementations based on WLAN are the most cost effective for the level of accuracy reached. Moreover, it can be combined with other methods using a different medium to increase the robustness.

B) **Cellular Network:** Accurate GSM-based indoor localization is possible as a result of wide signal-strength fingerprints that sense the six strongest GSM cells and up to 29 GSM channels which normally are too weak for communicating but strong enough to be detected and localized. The cells-ID are then registered as a fingerprint and treated in the standard way. For this method to work, the area should be covered by a GSM/CDMA cellular network. The accuracy depends on the network density: the denser a network is, the more accurate the localization will be. Experiments show that in densely covered areas the precision of GSM-based localization reaches 2.5 meters[14]. It is possible to exploit the public phone network, but a proprietary setup based on GSM is not convenient, due to the high cost of a network cell.

C) **Bluetooth:** With latest versions of Bluetooth it is possible to query the RSSI from known Bluetooth stations, along with other control values such as link quality and transmitted power control, and calculate the distance from the device in order to localize. Applying triangulation and a correction algorithm called Kalman filter, it was possible to increase the accuracy to 2.11 meters[15]. Recent improvements in Bluetooth technology are pushing towards a better robustness while decreasing the cost for the transmitters.

D) **Magnetic field:** Experiments[4] show that is possible to use the magnetic field as a fingerprint inside a building. This method is considered one of the most reliable because the magnetic field of a building does not change over time. The magnetic field is influenced by the structure of the building and it varies enough to offer a discrete precision and accuracy. It is even possible to locate a device in overlapping corridors on different floors. The minimum accuracy was 4.7 meters, but in the 50% of the experiments the target was localized with 2.5 meters of accuracy. The authors found that high power electrical machinery and devices that produces electromagnetic waves, such as laptops and smartphones, could cause interference with the real magnetic field and introduce errors in the readings. Fortunately, the error drops into an acceptable range at 12cm from these devices.

2.2 Algorithms

Once the fingerprints are collected and organized in a database after the offline phase, there is the need for pattern-recognition algorithms for identifying the fingerprint read by the user of the service having the correct location stored in the database. This is the most delicate part of the whole system, because of the ambiguous nature of the available data. The first problem is that the offline map is made of discrete points, while the user is moving in a continuous space; the algorithms should be able to correctly interpolate between the reference points on the map and perform an intelligent guess in the case that the same fingerprint figures more than once in the database. The following list explains the most common algorithms used for fingerprint indoor localization.

A) **K Nearest Neighbour:** or simply K-NN, is an algorithm that influences the construction of the offline map to reduce the errors in the online phase. The area to analyse is first subdivided into a rectangular grid of

"candidate blocks"; in each block RSSI samples are collected at a desired resolution. Every sample collected is then marked with the coordinate of the block (x, y) and the list of RSSIs/IDs present at that specific point. During the online phase, the user's device collects an array of the RSSIs signals sensed in its location and requests the base station to compare the array with the stored values in the database. The algorithm will find the *K* closest matched points of known locations in signal space from the previously-built database and return the average of their coordinates to the user's current location. The K-NN algorithm is preferred due to its easy implementation and relative high performances[9].

B) **Probabilistic Method:** While the deterministic algorithms, such as the kNN, estimate the position only by considering measurements of RSSs or the average and/or the variance of these values, the probabilistic algorithms compute the position by considering measurements as part of a random process, trying to exploit all the information contained within the signals acquired. It models the location fingerprint with conditional probability and utilizes the Bayesian inference concept to estimate the user position, together with other algorithms such as Nearest Neighbour[3]. During the online phase, the user's device to localize collects the RSSIs from the environment and the base station calculates the probability of the user to being at any of the reference points in the map. The rule for choosing the candidate node N_i is: choose N_i if $P(s|N_i) > P(s|N_j) \forall\ i,\ j = 1, 2, 3 \ldots n$ with $j \neq i$[10]. It means that, for all the recorded points j in the radio map, calculate the probability that the user is at the point i by using a Gaussian distribution, then choose the point with the highest probability or use K-NN with the K points of highest probability. The probabilistic method results accurate in most cases reaching 1.6 meters of accuracy.

C) **Neural Networks:** The database collected in the offline stage can be used for the training of a neural network that takes a number of RSS/APNs as input and returns a vector of two or three elements, which are the user's coordinates in 2D or 3D respectively [7].

D) **Cluster filtered K-NN:** Cluster Filtered K-NN (CFK)[11] is a technique that enhances the standard K-NN algorithm using a clustering technique to partition neighbouring points into sets to filter some of the matching neighbours to increase the accuracy. There are two kinds of clustering: **agglomerative** and **divisive**. The first method calculates the K-NN neighbours of a fingerprint, creates as many groups as the resulting neighbours and then merges the closest groups two by two; if the minimum distance d_{min} for each pair of groups is greater than a defined threshold *T*, stops the partition; otherwise, proceeds with another merge stage. The second method does the inverse: it starts with one big group which includes all the neighbouring points and splits it into two until the distance condition is met. At the end of the partition, there will be a handful of potential clusters in which the exact location is contained. The next step is applying one or more ***rule set R*** to select the candidate cluster, for example "select the cluster with more elements" and "select the cluster with smaller average RSS distance". Once the final cluster is selected, the coordinates are taken as the average of all samples in the set. Experiments[11] show that CFK reduces the prediction error of the normal K-NN algorithm.

3 CASE STUDY: proposal of an indoor localization application

This section shows the use cases and the architectural overview of an experimental application for indoor localization service using fingerprints. The application will be developed for smartphones and available for Aalto students as beta testers, the building of reference for the test will be the Aalto CS building in Konemiehentie 2. The application should fulfill the following requirements:

- The client device (smartphone) should be able to read the WiFi APNs and RSSI (Radio Signal Strength Indicator) from the environment using the internal wireless adapter.
- The smartphone should read the magnetic field from the environment using the internal magnetometer.
- The smartphone should be used to collect data from the building to create the radio map, and send the data to a server in real time or later.
- The smartphone should establish a continuous connection with the server during the online stage and query the location at defined intervals of time.
- A server should accept raw radio data from a client device and store it in a database organized in meaningful fingerprints.
- A server should run a pattern recognition algorithm to infer the client's position either in coordinates (x, y) or with its corresponding label ("Room A123") and send it back.

A scheme for the application architecture is shown in Fig. 1, separated in two distinct server and client components.

The server is agnostic about the client devices and its only task is to receive and elaborate data from the devices that request the service. The client have two main operational modes: online and offline. An example of the user interface is shown in Fig. 3. In the offline mode the client can work without a communication with the server, it should store the collected samples in a temporary storage and then synchronize with the server once the collection is done. To facilitate the construction of the radio map, the user will assign manually a unique name to the samples recorded. For example, if the user wants to collect data samples from the room A123, she must write "Room A123" in the UI text field before pressing the button tu record a sample. Then, all the sample collected will be saved with the given reference.

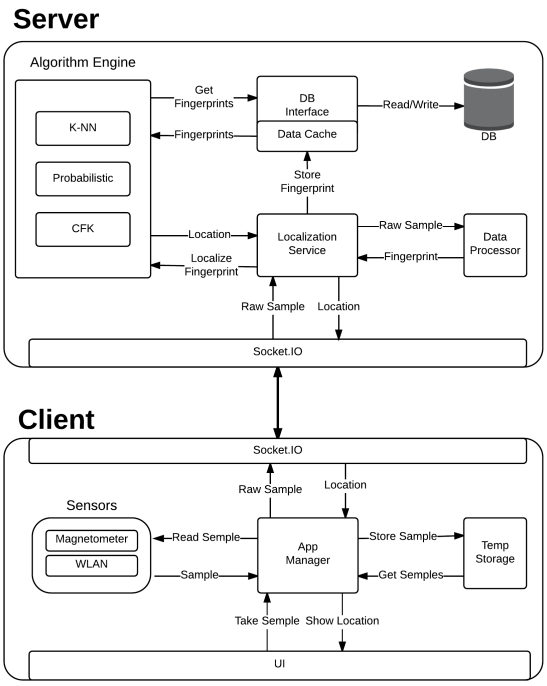

Figure 1: High level architecture

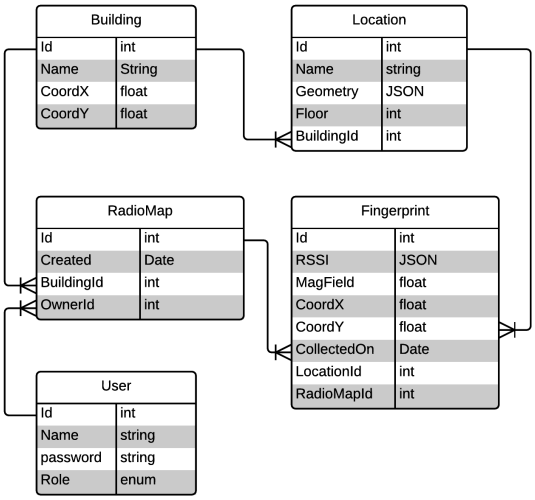

Figure 2: Fingerprints Database

The Figure 2 shows the proposed database structure for saving the fingerprints. Firstly, the localization service needs to organize the collected fingerprints in a way that allows different levels of precision inside the same building, at the same time the service should be able to save information relative to multiple buildings. The **Building** entity stores the world X and Y coordinates of the building so there could be an integration with existing map services. A **Location** is a 2D representation of a space in a map: it can be a room, a corridor, an hall and anything else that needs to be described on a map; it has a floor number and a geometry description, saved in GeoJSON format[1], an open standard used for encoding a variety of geographic data structures. There are tools for creating, exporting and visualizing maps that support GeoJSON, thus facilitating the process of creating the maps of the buildings. A location may have one or multiple **Fingerprints** associated, containing all RSSI values and IDs registered at a specific location in a JSON array. A fingerprint also includes information of the magnetic field and, if available, the precise X and Y position of the point in the map where the fingerprint is collected. With this organization, the application can localize the position of a user in three different precision levels (on a floor, inside a room, at a precise point). A **Radiomap** is a set of fingerprints collected in the building, kept separated for convenience and testing purposes. The server can switch radiomap to run the pattern matching algorithm and try different approaches using different configurations of fingerprints collected from the environment. In addition, there can be multiple **Users** that create radiomaps, and is important to prevent the pollution of the optimized radiomaps edited by the researchers. For protecting the user's privacy and grant security, he/she need to register and be logged in order to use the service. Moreover, a role subdivision mechanism assign grades to the users and restricts their actions accordingly.

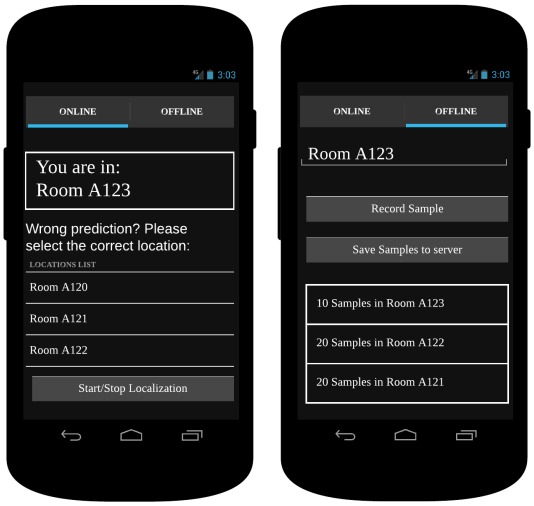

Figure 3: Application's UI mockup

In this experimental application both RSSIs/APNs and magnetic field information are collected in order to increase the accuracy of the localization using only WiFi information. For flexibility reasons the client will not perform any pre data processing, so the server can be changed to adopt different data analysis algorithms without having to change the client software. After that the data collection

phase is finished, the client can send all the information collected at once to save energy. Before the beginning of the online phase, the server should perform a preliminary data analysis in order to check the validity of the data, optimize it for being stored and eventually run a filtering algorithm to exclude points which may compromise the accuracy of the prediction. In [6] is shown a technique to discard fingerprints and APs from the samples, increase the overall application performance following simple rules. For example an AP which has a bad distribution over the radio should be discarded, or if it has a high percentage of missing RSSI values compared to other APs.

After the filtering, the raw samples collected from the client will be converted in fingerprints identified by their label and stored in the database. Since all database operations require a call to the file system, the fingerprints are chached in memory for a faster access. The client initiates the online phase opening a communication with the server using the TCP protocol for a fast and continuous stream of information. During the online phase, the client will collect raw samples of RSSIs/APNs and magnetic field information from the environment and send them to the server. At this point, the server can run one pattern recognition algorithm of choice (K-NN, Probabilistic or CFK), which is useful for experimental purposes and for collecting data about the performance of different algorithms. After the elaboration, the algorithm engine should have recognized the position of the fingerprint recorded by the client, and the server can send it back. If the prediction happens to be wrong, the user can correct it by selecting the correct position from a proposed list of locations, the server should register the correct information and update its prediction algorithm to take the correction into consideration. A simple neural network can be trained by means of user feedbacks and help the pattern recognition algorithms to calculate the exact location.

4 MAIN CHALLENGES

This section describes the challenges for a typical localization application using fingerprinting. A common issue addressed many times in the literature is the long and tedious process of collecting the first samples for creating the radio map of the building. Another important issue that affects the indoor localization in general is the error mitigation and error prevention: a wireless signal is not fully reliable due to its physical nature, thus the algorithms should try to actively reduce the error while the clients are using the service. One method is to run a filter to exclude unnecessary fingerprints from the radio map, as explained in Section 3; a second method could be to have direct feedback from prediction errors from the users of the service, and then train a neural network to recognize and correct the result in real time. The most part of the effort should be put in the development of pattern-recognition algorithms aimed to minimize the errors in predicting the location.

To facilitate the work of collecting fingerprints for the radio map, the data can be crowd-sourced with the help of Aalto students. For this purpose, the application could include gamification elements to encourage the students to participate in the experiment and involve even more participants.

Security and privacy are important issues for IPSs, since they focus mainly on user's activity. At any moment the user should be able to decide when to start and stop the tracking service, and the system should prevent uncontrolled access to the user's personal data, such as location history. The enhancements of security and privacy could be carried out from the software side: for example, by establishing a secure connection with the server, encrypt user data and grant access only after a login process. A self-localized position system architecture, instead, can ensure the privacy by performing location estimations directly in the target device if it knows the whole radio map. Unless the target device gives its location information to an entity, no one can access the information. Thus IPSs with self-localized location computation architecture can offer a high degree of security and privacy for the users, even if other problems rise, for example the computational power required to run the pattern matching algorithm in the device and the storage space needed to contain the radio map of the building.

5 CONCLUSION

This paper provides a clear survey of different techniques for indoor positioning systems using fingerprinting, comparing them in terms of cost, accuracy, robustness and reliability. Several sensing mediums such as WLAN, bluetooth, magnetic fiel, are explained in section 2 highlighting the differences between their characteristics and drawbacks; finally there is a discussion about the most commonly used algorithms for inferring the user's location such as K Nearest Neighbour, probability method and clustering. The paper proposes and describes a simple demo application for localizing a user inside the Aalto CS building that implements an indoor positioning system collecting WLAN and magnetic fingerprints from the environment, useful for testing different combination of algorithms and evaluate their performances. The application will be further developed as part of a thesis project as a tool for testing new algorithms and indoor localization methods.

References

[1] GeoJSON open standard. `http://http://geojson.org/`.

[2] A. Bekkali, H. Sanson, and M. Matsumoto. Rfid indoor positioning based on probabilistic rfid map and kalman filtering. In *Wireless and Mobile Computing, Networking and Communications, 2007. WiMOB 2007. Third IEEE International Conference on*, pages 21–21. IEEE, 2007.

[3] I. Bisio, F. Lavagetto, M. Marchese, and A. Sciarrone. Performance comparison of a probabilistic fingerprint-based indoor positioning system over different smartphones. In *Performance Evaluation of Computer and*

Telecommunication Systems (SPECTS), 2013 International Symposium on, pages 161–166. IEEE, 2013.

[4] J. Chung, M. Donahoe, C. Schmandt, I.-J. Kim, P. Razavai, and M. Wiseman. Indoor location sensing using geo-magnetism. In *Proceedings of the 9th international conference on Mobile systems, applications, and services*, pages 141–154. ACM, 2011.

[5] T. L. D. Yun, H. Chang. Accelerating vision-based 3d indoor localization by distributing image processing over space and time. In *ACM Virtual Reality Software and Technology*, 2014.

[6] S. Eisa, J. Peixoto, F. Meneses, and A. Moreira. Removing useless aps and fingerprints from wifi indoor positioning radio maps. In *Indoor Positioning and Indoor Navigation (IPIN), 2013 International Conference on*, pages 1–7. IEEE, 2013.

[7] L. Gogolak, S. Pletl, and D. Kukolj. Indoor fingerprint localization in wsn environment based on neural network. In *Intelligent Systems and Informatics (SISY), 2011 IEEE 9th International Symposium on*, pages 293–296. IEEE, 2011.

[8] Y. Gu, A. Lo, and I. Niemegeers. A survey of indoor positioning systems for wireless personal networks. *Communications Surveys & Tutorials, IEEE*, 11(1):13–32, 2009.

[9] T.-N. Lin and P.-C. Lin. Performance comparison of indoor positioning techniques based on location fingerprinting in wireless networks. In *Wireless Networks, Communications and Mobile Computing, 2005 International Conference on*, volume 2, pages 1569–1574. IEEE, 2005.

[10] H. Liu, H. Darabi, P. Banerjee, and J. Liu. Survey of wireless indoor positioning techniques and systems. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 37(6):1067–1080, 2007.

[11] J. Ma, X. Li, X. Tao, and J. Lu. Cluster filtered knn: A wlan-based indoor positioning scheme. In *World of Wireless, Mobile and Multimedia Networks, 2008. WoWMoM 2008. 2008 International Symposium on a*, pages 1–8. IEEE, 2008.

[12] E. Martin, O. Vinyals, G. Friedland, and R. Bajcsy. Precise indoor localization using smart phones. In *Proceedings of the international conference on Multimedia*, pages 787–790. ACM, 2010.

[13] D. Navarro and G. Benet. Magnetic map building for mobile robot localization purpose. In *Emerging Technologies & Factory Automation, 2009. ETFA 2009. IEEE Conference on*, pages 1–4. IEEE, 2009.

[14] V. Otsason, A. Varshavsky, A. LaMarca, and E. De Lara. Accurate gsm indoor localization. In *UbiComp 2005: Ubiquitous Computing*, pages 141–158. Springer, 2005.

[15] F. Subhan, H. Hasbullah, A. Rozyyev, and S. T. Bakhsh. Indoor positioning in bluetooth networks using fingerprinting and lateration approach. In *Information Science and Applications (ICISA), 2011 International Conference on*, pages 1–9. IEEE, 2011.

Big Data Platforms Supporting SQL

Markku Hinkka
Student number: 44157B
Markku.Hinkka@aalto.fi

Abstract

This paper surveys the most widely used Structured Query Language (SQL) supporting Big Data platforms. It discusses their design principles and identifies their possible strengths and weaknesses. Platforms are also categorized and compared against each other. Paper also includes a brief look into performing the selection of a platform and discusses the current trends and the future of SQL Big Data platforms.

KEYWORDS: SQL, Platform, Big Data, Hadoop, Dremel, Spark, Presto, Impala, Hive, Tajo, Drill, Amazon RedShift, Pivotal HD:HAWQ, Oracle Big Data SQL, IBM BigSQL, Teradata (SQL-H), Microsoft Azure HDInsight, Google BigQuery, Google F1

1 Introduction

In the last ten years, the amount of data stored into databases and data warehouses has increased dramatically. Companies' operations rely more and more on systems that generate all kinds of logs and other new data. This data is stored into their own data warehouses for analysis and to be used, for example, to monitor as-is-processes of the company and to detect bottlenecks. As a result, every 2 days we create as much information as we did from the dawn of man up to the year 2003 [8]. Storing and managing this ever growing mountain of data becomes harder and harder for traditional platforms that were not designed for extreme scalability [6].

Analyzing these vast amounts of data is usually very expensive, time consuming and often even impossible using the more traditional database management systems. However, in the last few years, various new platforms have been published that allow making complex analysis even of very large data sets.

This paper investigates the approaches and designs of some of the most widely used Big Data platforms and compares their differences with each other. It focuses only on platforms used to store structured data, thus NoSQL platforms are excluded.

The rest of this paper is organized as follows: Section 2 focuses on describing the meaning of Big Data as a term and also the problems that Big Data platforms help to solve. Section 3 discusses various most widely used platforms, their implementation principles and architecture. Section 4 presents a tabular representation of open source platforms and their differences. Section 5 explains briefly how the selection of a platform could be performed. Finally Section 6 explores the direction in which Big Data platforms seem to be moving at the time of writing. This includes considerations of both those platforms that are gaining popularity and those that are declining in popularity. This section also explores whether there are any new technologies that have potential to become the next "Big Thing" in Big Data.

2 Big Data

Big Data is a term that has been gaining more and more popularity lately. Usually it refers to large amounts of data that are extremely hard to analyze using traditional methods. Adam Jacobs in his article [6] suggests the following definition for Big Data at any point in time as: "data whose size forces us to look beyond the tried-and-true methods that are prevalent at that time". A HACE Theorem [14] describes the Big Data characteristics as follows: "Big Data starts with large-volume, heterogeneous, autonomous sources with distributed and decentralized control, and seeks to explore complex and evolving relationships among data".

Big data often consists of data that has been copied from companies' operative systems to data warehouses for reporting and analysis purposes. The data in data warehouses is also often denormalized, resulting in it taking more resources to store the data than it did in the original operative system [9]. The purpose for denormalization is to optimize the analysis tasks performed on the data.

Since traditional methods are no longer effective, new Big Data platforms have been developed to cope with these large amounts of data. Big Data platforms usually provide means to both store and analyze large amounts of data efficiently by distributing load evenly to varying numbers of worker hosts.

3 Platforms

One of the first and one of the most wide-known such platforms is Apache Hadoop [5], which is based on the MapReduce computing framework [4] originally developed at Google for generating the indexes needed for their Web search engine. The Hadoop implementation is used by many companies, including Yahoo and Facebook, for processing data in their massive data warehouses.

Apache Hadoop and many similar platforms include, in one form or another, four modules: Coordination for scheduling and resource management, processing for parallel processing, file system for accessing application data, and support utilities. In Apache Hadoop, parallel processing has been historically mostly performed using MapReduce [4] algorithm. Whether an operation can be parallelized efficiently

using MapReduce algorithm depends heavily on the upper limit of the number of inputs a reducer may receive [12]. In other terms, this can be generalized by saying that the more source data that is required to generate one result data item, the fewer possibilities there are for parallelization.

Big Data platforms can be divided into two categories based on the type of the data it supports: Structured and unstructured. Structured platform provides the means for storing and accessing data in similar manner to traditional relational database management systems (RDBMS) having similar constructs such as tables, rows etc. Accessing data in structured platform is usually performed using language similar to Structured Query Language (SQL). In unstructured platforms, data is usually not stored as tables having rows and columns. Instead, there is no predefined data model the stored data must follow.

A study by Y. Chen et al. [1] indicates that, at least for the selected few business-critical deployments examined in the study, there were only a handful of different analysis platforms used, of which quite a big fraction of tasks were run using structured platforms such as Apache Hive. It also seems that structured platforms make it easier for enterprises to transition from traditional RDBMS based platforms to structured SQL supporting Big Data platforms, since algorithms can be migrated often almost without any changes required. Enterprise users are often also more familiar with SQL queries than, e.g., MapReduce algorithms written in Java.

This paper compares a few selected SQL supporting Big Data platforms. The selection was mostly based on their popularity. Some interesting platforms were also scoped out due to platforms' immaturity (e.g., BlinkDB, which has only developer alpha available). The following subsections introduce each platform, describe some potential motivators for their use. Finally, a comparison of each open source platform is presented in tabular format listing the most distinctive properties of the selected platforms. Popularity of a platform is measured by querying the number of matching web pages for the name of the platform at the time of writing.

Platforms are categorized based on their licenses: Platforms having open source license and platforms having proprietary license.

3.1 Open Source Platforms

Traditionally the most popular Big Data platforms (e.g., Hadoop) have been released under free or open-source software licenses. The following subsections give a brief summary of the design guidelines of some of the most popular SQL supporting open-source platforms. For the selected platforms, all are released under Apache 2.0 license[1].

3.1.1 Apache Hive

Being one of the first distributed platforms supporting SQL, selecting Hive was quite natural. Hive is considered one of the de facto tools installed on almost all Hadoop installations [2]. It is still very actively being developed (1.0 version was released on February 2015) and has probably the biggest feature set of all the selected platforms. Hive was originally developed for processing long-running data-heavy batch jobs, thus it is not really well suited for interactive use. Hive works on top of an existing Hadoop cluster. It transforms all the SQL queries given to it into one or many MapReduce-tasks that are passed to Hadoop for final processing [10]. Every task having an inherent latency originating from the overhead associated with using Hadoop's MapReduce services [3] makes interactive usage problematic.

Hive is best used for summarizing, querying, transforming and analyzing large sets of structured data. It is usually not the best choice for queries requiring fast response and especially not for interactive use [4].

3.1.2 Apache Spark SQL

Apache Spark claims to be "a fast and general engine for large-scale data processing" that can "run programs up to 100x faster than Hadoop MapReduce in memory, or 10x faster on disk" [5]. It is based in Resilient Distributed Datasets (RDD), which is a fault-tolerant collection of elements partitioned across the nodes of cluster that can be operated on in parallel [15][16]. By default, spark programs are implemented using e.g. Java, Scala or Python programming language. Spark itself can be run in standalone mode or, e.g., on Hadoop cluster. Spark SQL is the second SQL on Spark platform with Shark being the first. Shark has been discontinued in favor of Spark SQL and Hive on Spark. [6]

Spark SQL is an additional programming library on top of Spark's own libraries which allows mixing SQL queries into Spark programs. It claims to allow user to run unmodified Hive queries on existing Hive warehouse [7], but in practice a few non-supported features are still missing. Based on the on-line documentation [8], it is hard to say whether these Hive functionalities are available also without Hive backend. However, it would seem that the current implementation is still less than the one that is claimed [9] [10].

Spark, like most in-memory data set based solutions, is best suited for situations where multiple different analyses need to be performed on the same data set. In these kind of situations, performing the first analysis will load the data set from the disk into Spark's in-memory RDDs. All the subsequent analyses will be performed using these in-memory RDDs, making them finish much faster than if the data had been read from the disk for every analysis.

[1] https://www.apache.org/licenses/LICENSE-2.0

[2] http://blog.matthewrathbone.com/2014/06/08/sql-engines-for-hadoop.html

[3] https://amplab.cs.berkeley.edu/benchmark/

[4] http://opensource.com/business/15/3/three-open-source-projects-transform-hadoop

[5] https://spark.apache.org/

[6] https://phdata.io/the-truth-about-sql-on-hadoop-part-2/

[7] https://spark.apache.org/sql/

[8] http://spark.apache.org/docs/1.2.0/sql-programming-guide.html

[9] http://www.aproint.com/aggregation-with-spark-sql/

[10] http://apache-spark-user-list.1001560.n3.nabble.com/SQL-Is-RANK-function-supposed-to-work-in-SparkSQL-1-1-0-td16909.html

3.1.3 Apache Drill

Apache Drill claims to be "a low latency distributed query engine for large-scale datasets, including structured and semi-structured/nested data" [11]. In the same way as Cloudera Impala, which is explained later in the paper, it is inspired by Google's Dremel. Dremel is an interactive ad-hoc query system for analysis of read-only nested data stored in a columnar storage [7]. Queries are executed using multi-level execution trees, making it capable of running aggregation queries efficiently.

It is recommended that the Drill is installed into an already existing Hadoop cluster. However, the minimum requirement is that it requires Apache Zookeeper [12] to be installed into the cluster. Drill can also be used to query some NoSQL databases such as MongoDB [13].

According to the documentation, Drill supports the ANSI standard for SQL 2003. It has some extended functionalities available for accessing nested and more complex data stored into table columns using map and array data types. It has several extensibility points such as pluggable query languages as well as support for user-defined functions (UDF). Drill's optimizers can be extended and support for new data sources and file formats can be implemented via API. At the time of writing, drill is still in beta phase of development (version number 0.7).

Drill is best suited for low latency queries, especially on semi-structured data or data stored in supported NoSQL databases. It is also a useful solution when there are no other platforms supporting the data source types that are needed.

3.1.4 Apache Tajo

Apache Tajo claims to be "designed for low-latency and scalable ad-hoc queries, online aggregation, and ETL (extract-transform-load process) on large-data sets stored on HDFS (Hadoop Distributed File System) and other data sources" [14]. It runs on top of Hadoop cluster. The online documentation is lacking and it seems that user base is not very big outside South Korea. It is mostly being developed by a South Korean startup called Gruter [15].

Tajo's is ANSI SQL standard compliant. It has support for some more complex SQL functionalities such as windowing functions, regular expression support for strings and it also supports GeoIP [16] functions out of the box. Tajo also provides dynamic load balancing and fault-tolerance for long running queries.

Tajo is probably best suited for performing low-latency ETL related tasks on data stored into HDFS. Its set of supported SQL functions seem to be quite comprehensive especially for ETL purposes. On the other hand, the seemingly small user base and poor documentation raises some questions regarding platform's future.

3.1.5 Facebook Presto

Facebook Presto claims to be "an open source distributed SQL query engine for running interactive analytic queries against data sources of all sizes ranging from gigabytes to petabytes" [17]. Presto was designed to handle data warehousing and analytics: data analysis, aggregating large amounts of data and producing reports [18].

It allows querying data from various sources such as Hive and Cassandra and even more traditional relational databases such as MySQL and PostgreSQL. It also provides an interface for building new custom connectors to different kinds of data sources. Presto was initially created by Facebook and it is still undergoing development by Facebook internal developers and a number of third party developers in Presto's community.

Presto employs a custom query and execution engine with operators designed to support SQL semantics. All processing is performed in-memory and pipelined across the network between stages, thus avoiding unnecessary I/O and associated latency overhead [11]. Presto supports standard ANSI SQL including extensive set of data types and functions including windowing functions and approximation based aggregation functions. Current version does not support UPDATE-operations; INSERT, however, is supported.

Presto is best suited for making quick interactive relatively simple queries on supported data sources or when a custom data source needs to be used. I personally have been having problems with some more complex queries in Presto [19]. Also it seems that if the query processing worker host runs out of memory, there is really no fallback to, e.g., using disk storage, and the query will just fail.

3.1.6 Cloudera Impala

Cloudera Impala claims to be "a fully integrated, state-of-the-art analytic database architected specifically to leverage the flexibility and scalability strengths of Hadoop - combining the familiar SQL support and multi-user performance of a traditional analytic database with the rock-solid foundation of open source Apache Hadoop and the production-grade security and management extensions of Cloudera Enterprise." [20].

Like Apache Drill, Cloudera Impala is inspired by the ideas in Google's Dremel [7] technology. Impala utilizes the same meta data, ODBC driver, SQL syntax and user interface as Hive. It is designed to complement the use of Apache Hive when interactive use is desired. It employs its own massively parallel processing (MPP) architecture on top of HDFS, thus it requires Hadoop to be deployed into the same cluster. It is also an integrated part of a Cloudera distribution of Apache

[11] https://cwiki.apache.org/confluence/display/DRILL/Architectural+Overview
[12] http://zookeeper.apache.org/
[13] https://cwiki.apache.org/confluence/display/DRILL/MongoDB+Plugin+for+Apache+Drill
[14] http://tajo.apache.org/
[15] http://siliconangle.com/blog/2015/03/11/apache-tajos-big-data-warehouse-comes-to-hadoop/
[16] http://dev.maxmind.com/geoip/legacy/downloadable/
[17] https://prestodb.io/
[18] https://prestodb.io/docs/current/overview/use-cases.html
[19] https://groups.google.com/d/msg/presto-users/5BI9Pb0mGN0/6ioSQn9O9SYJ
[20] http://www.cloudera.com/content/cloudera/en/products-and-services/cdh/impala.html

Hadoop (CDH) that includes various tools bound into a single package providing, among others, a common user interface and access control. In addition to Impala, it includes, e.g., Hadoop, Hive, Pig and Spark. [21]. Cloudera is the only contributor and developer to Impala code [22].

Impala is best suited for interactive SQL analysis cases especially when commercial support is required or when other tools included in CDH are required to be installed into the same cluster, e.g., for stream processing purposes. According to a tests performed by Wouw et al. [13] Impala seems to be best suited for input sizes smaller than 500 GiB. However, there seem to be also somewhat contradicting slightly older results in similar performance tests [2].

3.2 Commercial Platforms

There are quite a few commercial platforms, licensed under proprietary licenses, competing with open source platforms. The following subsections give a brief summary of some of the most popular ones.

3.2.1 Amazon RedShift

Amazon RedShift claims to be "Fast, fully managed, petabyte-scale data warehouse solution that makes it simple and cost-effective to efficiently analyze all your data using your existing business intelligence tools." [23]. It is Amazon's hosted data warehouse product, which is part of Amazon Web Services computing platform. Redshift is based on massive parallel processing (MPP) data warehouse ParAccel [24]. Its Performance and capacity can be elastically scaled on-demand without downtime, it supports data updates and transactions and has pricing based on the actual usage. It performed consistently well in AmpLab's Big Data benchmark [25].

3.2.2 Pivotal HD:HAWQ

Pivotal's HD:HAWQ claims to be "a parallel SQL query engine that combines the key technological advantages of the industry-leading Pivotal Analytic Database with the scalability and convenience of Hadoop." [26]. HAWQ is part of Pivotal HD, which is Pivotal's proprietary Hadoop distribution. It runs on top of Hadoop cluster that uses Pivotal's Hadoop distribution. It is based on Greenplum database with modifications that enable data to be stored into HDFS. It has MPP SQL processing engine optimized for analytics with full transaction support [27]. It is quite mature product having been under development for over 10 years [28]. It performed quite well on already quite old benchmark against Hive and Impala [29]. It also has full transaction support.

3.2.3 Oracle Big Data SQL

Oracle's Big Data SQL claims that it "Extends Oracle SQL to Hadoop and NoSQL and the security of Oracle Database to all your data." [30]. Big Data SQL uses a traditional database as front end for all the SQL queries. It requires both Oracle Exadata that contains traditional RDBMS database and Oracle Big Data Appliance that contains a Cloudera Hadoop cluster [31]. It supports multiple data sources, including Hadoop, NoSQL and Oracle Database. It enables users to write a single SQL query combining data from Oracle Database (versions 12c and up), Hadoop and NoSQL database.

3.2.4 IBM BigSQL

IBM BigSQL claims to be "Massively parallel processing (MPP) SQL engine that deploys directly on the physical Hadoop Distributed File System (HDFS) cluster" [32]. BigSQL is part of Big Insights, which is IBM's proprietary Hadoop distribution. It runs on top of an existing Hadoop cluster and uses Hive Metastore for data definitions. It does not require any proprietary Big Data Appliance and does not support transactions.

3.2.5 Teradata (SQL-H)

Teradata Enterprise Access to Hadoop claims to be "a portfolio of tools designed to give any Teradata user easy, secure access to data stored in Hadoop." [33]. Teradata runs on top of an existing Hadoop cluster, which does not have to be Teradata's own distribution. Teradata has also an own Hadoop distribution optimized for Teradata SQL-H. It enables users to issue a remote query to Hadoop and bring the data into Teradata traditional RDBMS for analysis or integration with the data already stored there. It does not require any proprietary Big Data Appliance.

3.2.6 Microsoft Azure HDInsight

Microsoft Azure HDInsight claims to be "A 100% Apache Hadoop service deployed in the cloud that can elastically scale out to big data within minutes" [34]. It is Microsoft Azure cloud implementation of Hadoop providing on-demand scaling. [35]. It uses Hortonworks Data Platform Hadoop distri-

[21] http://www.cloudera.com/content/cloudera/en/products-and-services/cdh.html
[22] http://www.slideshare.net/NicolasJMorales/big-sql-competitive-summary-aug-2014
[23] http://aws.amazon.com/redshift/
[24] http://www.zdnet.com/article/amazon-redshift-paraccel-in-costly-appliances-out/
[25] https://amplab.cs.berkeley.edu/benchmark/
[26] http://pivotal.io/big-data/pivotal-hd
[27] http://pivotalhd.docs.pivotal.io/doc/2010/HAWQOverview.html
[28] http://blog.matthewrathbone.com/2014/06/08/sql-engines-for-hadoop.html
[29] http://www.dataintoresults.com/2013/09/big-data-benchmark-impala-vs-hawq-vs-hive/
[30] http://www.oracle.com/us/products/database/big-data-sql/overview/index.html
[31] http://www.slideshare.net/NicolasJMorales/big-sql-competitive-summary-aug-2014
[32] http://www-01.ibm.com/support/knowledgecenter/SSPT3X_3.0.0/com.ibm.swg.im.infosphere.biginsights.product.doc/doc/bi_sql_access.html
[33] http://www.teradata.com/Teradata-Enterprise-Access-for-Hadoop/
[34] download.microsoft.com/download/F/7/C/F7C2E119-A2EA-4660-8D8C-C6C55BB844EF/AzureHDInsightInfographic2015.pdf
[35] http://azure.microsoft.com/fi-fi/documentation/articles/hdinsight-hadoop-introduction/

bution and can be deployed using Linux or Windows as underlying OS. HDInsight includes, among others, an implementation of Hive which makes it possible to query data using HiveQL. Users have to pay only for the actual compute. It integrates with Microsoft's tools such as Excel, SQL Server Analysis Services and SQL Server Reporting Services. HDInsight makes it possible to develop MapReduce-jobs also using .NET languages.

3.2.7 Google BigQuery

Google's BigQuery claims to "Analyze Big Data in the cloud with BigQuery. Run fast, SQL-like queries against multi-terabyte datasets in seconds. Scalable and easy to use, BigQuery gives you real-time insights about your data." [36]. It is Infrastructure as a service hosted in Google Cloud Platform. It uses Google Storage as its data source. BigQuery is Google's own public implementation of Dremel [7] [37]. BigQuery charges for data storage and for querying data, but loading and exporting data are free of charge [38]. Also there is no need to pay for provisioning servers since the service is always on [39]. Google BigQuery also has a free trial period during which there are some product limitations.

3.2.8 Google F1

Google's F1 claims to be "a distributed relational database system built at Google to support the AdWords business. F1 is a hybrid database that combines high availability, the scalability of NoSQL systems like Bigtable, and the consistency and usability of traditional SQL databases." [40]. F1 is built on top of Spanner, which is a globally distributed NewSQL database and a successor of BigTable. Its key design goals are: Scalability, Availability, Consistency and Usability. Scalability in this context means that the platform must be able to scale up just by adding resources. Availability means that the system must never go down for any reason. Consistency means that the system must provide ACID transactions. Usability means that the system provides SQL query support and other functionality users expect from a SQL database. F1 seems to be only used internally by Google which uses it at least for its AdWords business.

4 Comparison

The Table 1 in the end of the paper lists some of the most distinctive properties of all the selected Open Source platforms shown as columns. The selected properties shown as rows are:

- **Community**: Developer community developing the platform.
- **Chief advocate**: Primary commercial advocate funding the development.
- **Required back end software**: Software that is required in the system in order to be able to use the platform.
- **SQL variant**: Supported SQL variant.
- **Fault tolerance**: Does the platform have built-in fault tolerance for queries or does the user have to handle faults.
- **Commercial support**: Does the platform have commercial support available.
- **Low latency**: Is the platform suitable for low latency and interactive queries.
- **In-memory**: Does the platform primarily rely on main memory for data storage?
- **Data shapes**: Supported shapes for the queried data.
- **Data sources**: Supported data sources.
- **ODBC/JDBC driver**: Does the platform have ODBC and JDBC drivers?
- **Supported data modifications**: What kind of data modifications are supported.
- **Sandbox available**: Prepared easy to use sandboxes available for evaluation purposes.
- **Google keyword**: Keyword used for estimating the popularity in Google matches-column.
- **Google matches**: Number of web pages matching the Google keyword in a search made using Google search.[41]

All the information in this seminar paper is based on the situation in April 2015. Since most of these platforms seem to be evolving very rapidly, some information given in this paper may already be obsolete when this seminar paper is published.

5 Selecting Suitable Platform

Since deploying a Big Data SQL platform into a cluster is usually quite complex task, it is very important to find out the best suiting platform for your requirements. This is often performed by weighting the pros and cons of the platforms against each other. Questions you could ask could include, e.g.:

- Are you ready to pay for the platform?
- Are you already familiar with some platform or the ecosystem of some platform vendor?
- Are you willing to dedicate a cluster of computers for this purpose? If not, do you want to pay for provisioning servers or only for the usage of actual computation resources?

[36] https://cloud.google.com/bigquery/

[37] https://cloud.google.com/files/BigQueryTechnicalWP.pdf

[38] https://cloud.google.com/bigquery/pricing

[39] http://www.quora.com/How-good-is-Googles-BigQuery-as-compared-to-Amazons-Red-Shift

[40] http://static.googleusercontent.com/media/research.google.com/fi//pubs/archive/41344.pdf

[41] https://www.google.com/

- Are you expecting that the same platform will be used also for some still more or less unknown future purposes?
- If commercial support is required, does the platform have it?
- Are you expecting to use the platform for a long time? Does the platform seem popular and good enough to be alive longer than you need it?
- Do you require fast response times and user interaction or could your workloads be processed as batch jobs? Do you need to be able to run multiple queries concurrently?
- Does the platform support a data source suitable for your needs?
- Does the platform have all the SQL query features that are required?
- Does the platform support the needed data shape and/or required operations?
- Do you already have a cluster that has to be used? Does it have some back end software that sets constraints for the selection of the platform?

It is often also a good idea to do platform evaluations first in small scale, e.g., using prepared virtual machines having operating stand alone version of the platform pre-installed (a.k.a. sandbox).

6 The Future

It would seem, based on the amount of resources and variety of available platforms, that SQL has already established its place as a De facto language for specifying analysis tasks also for Big Data. For this there are many reasons such as the ease of transforming legacy applications from traditional RDBMS into Big Data SQL platforms and the familiarity of SQL language to data analysts around the world.

Also based on the research done for this paper, it seems that the era of MapReduce is coming to its end, at least as the back end implementation of Big Data SQL platforms. From 14 products explored, only 2 were using MapReduce for processing SQL queries. MapReduce seems not to be optimal for processing SQL [48] especially when considering interactive usage [49].

Several competing paradigms such as Resilient Distributed Datasets (Spark) and Dremel (Drill, Impala) have proven their place especially for running quick interactive analyses on large datasets. Partly this speedup comes from the fact that these new paradigms are based on in-memory processing, where all the available system memory in every worker hosts is exploited to as great degree as possible. Processing analyses in-memory is usually much more efficient than processing from a shared storage such as disk. Also constantly decreasing memory prices enables improving cluster performance just by adding more memory into the worker hosts.

It remains to be seen what will eventually be the "Next MapReduce". However, it seems quite interesting that even in Cloudera, which is the chief advocate for Dremel based Impala, has added Spark into their Hadoop distribution. They have also been writing in favor of Spark [50].

Also since currently there are lots of small platforms having very similar sets of features, it is quite probable that some of them will not survive for long, especially if there is nothing that makes them standing out of the crowd.

Quite a few of the selected platforms have inter-query fault-tolerance built-in which allows building reliable systems from less reliable parts. In the future large companies will also want to be able to improve system responsiveness and to be able to build predictably responsive systems having small latencies out of less predictable parts [3]. In fact, there are also some platforms being developed, such as BlinkDB [51], that actually make it possible for users to trade-off query accuracy for the response time by employing some more complex dynamic sampling techniques.

[48] `http://vision.cloudera.com/impala-v-hive/`

[49] `http://blog.mikiobraun.de/2013/02/big-data-beyond-map-reduce-googles-papers.html`

[50] `http://vision.cloudera.com/mapreduce-spark/`

[51] `http://blinkdb.org/`

References

[1] Y. Chen, S. Alspaugh, and R. H. Katz. Interactive analytical processing in big data systems: A cross-industry study of mapreduce workloads. *PVLDB*, 5(12):1802–1813, 2012.

[2] Y. Chen, X. Qin, H. Bian, J. Chen, Z. Dong, X. Du, Y. Gao, D. Liu, J. Lu, and H. Zhang. A study of sql-on-hadoop systems. In *Big Data Benchmarks, Performance Optimization, and Emerging Hardware*, pages 154–166. Springer, 2014.

[3] J. Dean and L. A. Barroso. The tail at scale. *Communications of the ACM*, 56(2):74–80, 2013.

[4] J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.

[5] J. Dittrich and J.-A. Quiané-Ruiz. Efficient big data processing in hadoop mapreduce. *Proceedings of the VLDB Endowment*, 5(12):2014–2015, 2012.

[6] A. Jacobs. The pathologies of big data. *Communications of the ACM*, 52(8):36–44, 2009.

[7] S. Melnik, A. Gubarev, J. J. Long, G. Romer, S. Shivakumar, M. Tolton, and T. Vassilakis. Dremel: interactive analysis of web-scale datasets. *Proceedings of the VLDB Endowment*, 3(1-2):330–339, 2010.

[8] E. Schmidt. In *Techonomy conference*, August 2010.

[9] S. K. Shin and G. L. Sanders. Denormalization strategies for data retrieval from data warehouses. *Decision Support Systems*, 42(1):267–282, 2006.

[10] A. Thusoo, J. S. Sarma, N. Jain, Z. Shao, P. Chakka, S. Anthony, H. Liu, P. Wyckoff, and R. Murthy. Hive: A warehousing solution over a map-reduce framework. *Proc. VLDB Endow.*, 2(2):1626–1629, Aug. 2009.

[11] M. Traverso. Presto: Interacting with petabytes of data at facebook. 2013. [Online; accessed 12-March-2015].

[12] J. D. Ullman. Designing good mapreduce algorithms. *XRDS: Crossroads, The ACM Magazine for Students*, 19(1):30–34, 2012.

[13] S. v. Wouw, J. Viña, A. Iosup, and D. Epema. An empirical performance evaluation of distributed sql query engines. In *Proceedings of the 6th ACM/SPEC International Conference on Performance Engineering*, ICPE '15, pages 123–131, New York, NY, USA, 2015. ACM.

[14] X. Wu, X. Zhu, G.-Q. Wu, and W. Ding. Data mining with big data. *Knowledge and Data Engineering, IEEE Transactions on*, 26(1):97–107, 2014.

[15] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. J. Franklin, S. Shenker, and I. Stoica. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, pages 2–2. USENIX Association, 2012.

[16] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica. Spark: cluster computing with working sets. In *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing*, pages 10–10, 2010.

Feature	**Hive**	**Spark SQL**	**Drill**	**Tajo**	**Presto**	**Impala**
Community	Apache	Apache	Apache	Apache	Facebook	Cloudera
Chief advocate	Hortonworks	Databricks	MapR	Gruter	Facebook	Cloudera
Required backend software	Hadoop	None	Zookeeper	Hadoop	None	Hive Metastore
SQL Variant	HiveQL	Rudimentary ANSI SQL & HiveQL	Extensible, SQL 2003, ...	ANSI SQL	ANSI SQL	SQL/HiveQL subset
Fault tolerance	Yes	Yes	Yes	Yes	No	No
Commercial support	Yes	Yes	Yes	Yes	No	Yes
Low latency	No	Yes	Yes	Yes	Yes	Yes
In-memory	No	Yes	Yes	Yes	Yes	Yes
Supports running jobs concurrently	Yes[42]	Yes[43]	Yes[44]	No[45]	No?[46]	Yes[47]
Data shapes	Nested, tabular	RDD	Nested, tabular	Tabular	Tabular	Nested, tabular
Data sources	HDFS, HBase	Any data source supported by Hadoop	Extensible, e.g., HDFS, HBase, ...	Hive, HBase	Hive, Cassandra, MySQL, ...	HDFS, HBase
ODBC/JDBC driver	ODBC+JDBC	ODBC+JDBC	ODBC+JDBC	JDBC	JDBC	ODBC+JDBC
Supported data modifications	Transactions, Append, Update	Spark Streaming	None	Append	Append	Append
Sandbox available	Yes	Yes	Yes	No	No	Yes
Google keyword	apache hive	apache spark sql	apache drill	apache tajo	facebook presto	cloudera impala
Google matches	226,000	50,100	130,000	28,300	42,100	77,200

Table 1: Open source platform comparison matrix

Review of energy profiling methods for mobile devices

Antti-Iivari Kainulainen
Student number: 62833A
ajkainul@cc.hut.fi

Abstract

This paper reviews different parts of the mobile device energy estimation and different implementations of them.

KEYWORDS: mobile device, consumption, energy efficiency, consumption model, consumption estimation

1 Introduction

In recent years many kinds of powerful mobile devices have become very popular and common items and they are expected to be usable all the time. However, as the devices have transformed into powerful full featured computers with very resource demanding applications, the use times have reduced, which significantly diminshes their usefulness.

The devices are designed to be quite energy efficient but the problem arises from poorly designed energy hungry applications. For a long time, the problem has been that the energy efficient programming has been hard since there has been very few usable tools available for assessing energy efficiency. Fortunately, in recent years there has been more research about the subject and many tools and methods has been developed for measuring and estimating consumption.

The scientific publications usually propose a certain way of estimating consumption and usually a tool is presented which has been developed with the paper. Many of the papers does not present a whole new approach to the problem but usually improvements to a previous study. This is why there are a wide variety of papers available which concentrate on improving some feature but there seems to be very few papers about the overall picture of the whole process in general, which would help those who are new to the subject.

This papers tries to help in this problem. It presents an overall picture of the current state of the energy estimation and what has been done, why and how.

2 Background

There are many papers about the energy estimation of mobile devices. However, this paper concentrates on building a general overview of the whole estimation process and how different features have been implemented in different papers and why. Table 1 presentes all the tools which have been analyzed in this paper and how they differ from each other. All of the details presented in the table are covered separately in this paper.

2.1 Consumption estimation as a whole

It is good to note at the start that energy consumption estimation does not only cover the actual estimation but the whole process, which can be divided roughly into three different parts: consumption measurement, model building and estimation using a model.

The estimation using a model estimates the device consumption by monitoring the activities of device components[1] and then by using a model, which describes precisely how much components consume when performing different activities. The amount of energy consumed is estimated based on the extent they were used. The measurement of actual consumption is used to build this model for estimation.

Next these are examined in more detail by first examining the consumption measurement from hardware and after that how model is built from measured data and finally how this model is used to estimate consumption.

3 Hardware measurements

This section examines why most of the papers seem to base their modeling on measuring whole device consumption. After that, we look in more detail at different characteristics of the internal and external measurements.

3.1 Device vs. Component measurement

All the consumption models are, of course, based on knowledge of the actual consumption. The ideal approach would be if the consumption of each component could be measured in real time without any external hardware, which would not only make building models easier but also pointless because the consumption could be just measured directly. Unfortunately, none of the components currently are capable to do that. Usually, the only part able to measure consumption is the Battery Management Unit (BMU). The BMU monitors the batery temperature, voltage and the discharge current[5].

Another approach would be to attach external current meters to a circuit board of the device. This has been done at least with OpenMoko Freerunenr by A. Carroll and G. Heiser[2]. Even in this case the OpenMoko did have open hardware to ease the process, so this would be much harder on some other device which does not have hardware schematics available. Many modern device designs favor more System In Package or System On Chip (SIP or

[1] In this paper component refers to some subsystem on the device, like 3G modem, Wifi modem, CPU or display.

Table 1: The characteristics of different estimation programs.

Estimation Consumption model	Appscope[3](AS) Devscope[5](DS)	PoweTutor[12] PoweBooter	Eprof[7][8]	PowerProf[6]	Sesame[4]
Measurement of consumption					
Internal/external	Internal (DS)	Internal,External	Internal?	Internal	Itnernal
Measured value	Current (DS)	Voltage	???	Voltage/current	Current and voltage
Model construction					
Model type	System call, events (DS)	Utilization	System call	Api, Genetic algrithm	Utilization
Intrusion level	Userspace (DS)	Userspace only	kernel modification	Userspace	Implements new syscall
Online?	Yes (DS)	Yes	Yes	Yes	Yes
Generality	Unit specific (DS)	Unit, Battery dependent	???	Unit	Unit and use pattern
Training type	Excercising application (DS) Synchronized update (DS)	Excercise application	Excercising application Automatic FSM transition	Genetic algorithm excercise to optimize fitness function	User usage pattern
Building	Automatic (DS)	Automatic, Manual	Automatic	Automatic	Fully automatic, Validity checks
Estimation					
Model type	System call, events (AS)	Utilization (procfs)	Instrumented source code, System call logging	Api Application instrumentation	Utilization
Intrusion level	Kernel module (AS)	Userspace	Kernel modification	Userspace	Userspace
Granularity	Application (UID) (AS)	Whole system level	Routine, Thread, Application	Api calls to start/stop	Whole system level
Restrictions			Source code required	Source code required	

Note: Some papers does not neccessarily provide both estimation and modeling but if the other is implemented for evaluation in paper it is classified based on the tool
Note: AppScope uses model generated by DevScope which is why they are grouped together. AS and DS are used to differentiate the exact characteristics.

SOC) solutions. These compound chips does contain multiple chips in a same package[11], for example Bluetooth radio and WiFi radio inside same chip. This of course prevents from measuring the current drawn by one component. It is also quite obvious that this way would also require a laboratory environment and still would be really tedious to implement, making it impossible for the majority of mobile device users.

Of course it would be possible also to statically analyze the hardware and drivers to calculate exact consumption but since this would require deep understanding of hardware, access to closed source drivers and chip documentation, it would be infeasible if not impossible to conduct.

Finally, the only feasible choices for measuring the consumption would be to either use the internal voltage/current sensor on BMU or to replace the battery with an external power supply and measure the current drawn by the device externally. Both of these are valid options and both of them have great advantages compared to the other. Section 3.2 looks at them in more detail.

3.2 External measurement

External measurement is measuring the consumption of the device externally by adding current meter between the mobile device and battery or by replacing the battery with a power supply which is capable of measuring current drawn by the device. One such device is Monsoon Mobile Device Power Monitor[1] which is used in many papers for modeling or evaluating the model, for example Pathak et al.[7] used this device for model evaluation, Zhang et al.[12] used this device for initial model building for PowerTutor and Jung et al.[5] used this device to validate their results.

The internal current/voltage sensors usually have a low update rate and voltage sensors especially are quite inaccurate for measuring the consumption[5]. The same restriction do not naturally apply to selection of external measurement devices and the Monsoon device, for example, does give more accurate readings and faster, too.

However, there are disadvantages which makes these kind of measurements impractical at least from application development point of view. Firstly, the user must acquire the measurement device, which is usually expensive. Secondly the mobile device might have to be opened because many of the modern devices, like iPhones[9], do have an internal battery which cannot be replaced by the user. Thirdly, basic knowloedge of electronics is required. Lastly, because an external device is used, the online model building is not possible since another device is needed to conduct the measurements. From the point of view of the estimation tool these all together make this approach unfeasible for most of the end users, and also software developers who don't have access to the required hardware and do not have knowledge to use such devices.

However, even if the hardware measurements are impractical from the user point of view, it is worth noting that they do have an important role in evaluating research results. Most of the papers, which present estimation approach, such as the one which presents AppScope[3], also use hardware measurements to validate results. Since the external measurements can be very accurate and may have very high acquisition rate, it is a good reference point for comparison. Usually the estimations are evaluated by estimating consumption with it and comparing the estimated consumption to measured one.

3.3 Device internal measurements

More convenient approach is to use the internal BMU chip of a mobile device which is responsible of charging the battery and indicating the battery status. The actual consumed power is measured by the BMU by measuring the voltage or current drawn by the mobile device. Usually current sensors give accurate reading about the consumption. However, many of the mobile devices only have a voltage sensor which estimates the power from a voltage drop of the battery, which is much more inaccurate. Fortunately, most of the modern devices have current sensors but there might be still some which have not. For example, PowerBooter automatically builds a model by using a voltage sensor because of this[12].

Both of the sensors however suffer from low update rates, which makes it more difficult to identify the exact rela-

tionship between consumption and actions[5]. There are methods to overcome these limitations. These methods are closely linked to details of exercising program which is used to exercise hardware to extract different power characteristics. Because of this, these methods are covered in more detail in section 4.1.2, which describes different approaches to exercise the hardware..

With these improvements, internal measurements are one possible option for constructing a fine-grained and accurate power model. This is a very appealing option since it makes it possible to do the measurements without external hardware and build the model online inside the device and even completely without any user interactions.

Section 4 considers how to turn the measurement data into an estimation model which describes component consumptions and how to overcome these limitations caused by internal measurements.

4 Consumption Model

The most essential component in estimation power consumption is the consumption model. It links consumption to a measurable unit of work performed by hardware. The exact method varies between different approaches.

In practice, this means that a program will be run on the mobile device which monitors some certain functionalities of the hardware while measuring the current drawn by the device. Next, the acquired measurement data is processed and by using a method such as linear regression, the hardware activity states or events are linked to actual consumption. The model constructed this way is then used to estimate the consumption.

This section first looks in more detail at the consumption measurement from the model building point of view and after that how the measured data can be turned into an actual consumption model.

4.1 Measuring the Consumption

This section is about measuring the consumption while monitoring the component activity. This is a bit more complicated than it first sounds. There are two challenges: the inaccuracy of internal power sensor measurements and the complex consumption characteristics of the hardware. Section 4.1.1 investigates the power sensor related challenges more closely and after that the what kind of challenges the functional characteristics of the devices will produce.

4.1.1 Handling the BMU Inaccuracy

One of the biggest challenges, which is only related to internal measurements, is the inaccuracy of the measurement hardware. The sensors might give inaccurate readings and the update rates are low. Many of the papers present two possible solutions to this.

Both of these solutions are based on training the model with specific training program which does not only log the component activities and device consumption. It also activates different parts of the hardware based on schedule. The actual characteristics of these programs will be discussed more closely on next section but on this context it is sufficient to know that these programs does set the components into some specific state and then measure how much the device consumes.

The first method to overcome the low update rate of the sensor is to average the readings from a long period of time[4].

The second solution is to synchronize the state changes to the update rate of the BMU. This is the way DevScope handles this problem. It does automatically figure out the update rate and synchronize the test scenarios to BMU updates.[5]

Another problem, which is related to devices which does have only voltage sensor, is that the voltage of the battery does not actually tell anything about the consumption. Fortunately the consumption information can be extracted from this information. The voltage of Lithium ion batteries will change based on their charge level[10]. This way it is possible to calculate the difference of battery charge level which equals to consumed energy. However the readings are inaccurate, which can be handled like above, but another problem is that the discharge curve is battery specific and even worse the age and use does affect to discharge curve too which makes the curve even unit specific[12].

Of course the discharge curve could be measured with external measurement device but then the advantage gained by using internal measurement would be lost. Fortunately the discharge curve can be acquired for some battery models so the model can be generated for new units by using this information. Also the lack of discharge curve might not be a problem when considering the consumption. At least PowerBooter is capable of creating model which does not measure actual consumption but relative to battery[12]. This kind of approach could be sufficient because the exact current value wouldn't add much value to the results. Usually it is more important to see difference between different components, programs or functions.

4.1.2 Handling the Component Characteristics

Most of the consumption comes from activity of different components of mobile device, such as 2G/3G/4G modem, Wifi, Bluetooth, display, CPU and GPS. Most of these components does one specific task, which is not dependent from the others, so many of the components behavior is also almost independent so it is also convenient to measure these components activity separately. However the most notable exception to this are memory and CPU which are, of course, required to control the components.

On the component level the exact way of dealing with component depends on its type. Some characteristics may be determined on runtime but most the model of behavior is modeled based on prior knowledge about the device interal workings.

Most of the referenced papers, like Pathak et al. presenting Gprof[7], uses finite state machines to model most of the components consumption. This is good approach since the consumption of the component heavily correlates with power state in most cases. For example GPS does have almost static consumption when in use and zero consumption when inactive. Another good example is 3G modem which can be in

three different states: only listening for beacons, in slow and low power state and on fast high power state[12].

The exact transitions from one state to another does depend on many factors, like mobile device energy saving policy or network policies defined by operator on 3G[5]. For example device energy saving policy for GPS might govern that the GPS will be turned on only when location is requested and it will remain on for short time after each request to serve faster consequent location requests, for example from navigator software, while allowing the component to go sleep when it is no longer used.

When building a model the knowledge about the different power states can be used when building training program for the device. Most of the papers does the traning so that they will generate training schedule for the hardware so that they will go through all the different power states of components one at the time while measuring the consumption[4][5][12][7]. The exact characteristics does vary between different implementations.

It is also worth noting that the device also does have base consumption which comes from all the other parts of the device which are not directly related to any components.

4.2 Turning Measurements Into a Model

Regression is very common way of presenting the estimated consumption and it is used at least by PowerBooter[12] and DevScope[5]. It is a simple model for fitting predictive model to an odserved data. On 1, which is how it was described by Yoon et al.[3], the x_i represents vector of usage of component i and β_i is the power coefficient for component i. P_{base} is a base consumption of the device and P_ϵ is the noise which cannot be estimated by the model.

$$P = \sum_i (\beta_i \times x_i) + P_{base} + P_\epsilon \quad (1)$$

It is also possible to use non-linear regression model too. Non-linear model would also measure the dependency among components. However non-linear models does not significantly outperform the linear models. [3]

The regression models are not the only considerable choice for modeling the consumption. Kajergaard et al.[6] uses genetic model to construct the power model. However it requires heavy processing for the data, which is impractical to do online on the device so the actual model building has to be done offline or offloaded to server, like with PowerProf. This of course requires another computer or server which does the calculation. This does affect on usability but it is worth considering when assessing the improvements gained by using more complex algorithm.

5 Estimation with model

When the model has been built the actual estimation of consumption can be done. There are numerous approaches to actual estimation. As with model building there are no really single solution which would be ultimately best. Most of the approaches have some really good characteristic while they are lacking something else, like really fine-grained system call tracking requires low level access to the device while offering very useful results.

Both training the model and estimating with it does share some characteristics since both of them does monitor component states but one of them does monitor consumption and other tries to estimate it. However the estimation of the consumption is not as straightforward as building the model. When estimating the granularity of the data is very important factor. On this context the granularity refers to how accurately the component consumption startpoint, endpoint and owner can be determined.

First on two first sections we are going to look different approaches for monitoring the component activities to get fine-grained estimates and after that compare the effects of the implementation to the user experience.

5.1 Estimate granularity

The term granularity on consumption estimation is used to describe how accurately the consumption can be linked to functionality of the device. The granularity level depends on what kind of features are monitored by the estimation program, for example the program can poll CPU activity from /proc on linux based systems or the estimation program sets breakpoints to system calls to know when certain component is used. The exact estimate types are discussed in more detail on section 5.2.

It is also worth noting that the purpose of the estimation tool does also affect to choice of monitored features. Especially older papers, like Zhang et al. paper about PowerTutor[12], does concentrate on problem about estimation consumption in general while more recent papers, like Yoon et al. paper about AppScope[3], concentrates on refining the process and choosing features which does give most value to the user. For example the developer is very likely interested in function call level consumption while end user is more likely more interested how much does use of some specific application consume energy.

The exact granularity depends, of course, heavily on implementation and model type of the estimation program and there are many minor differences on their characteristics and same program or estimation framework can offer more than one option. However the granularity of the estimate can be simplified and classified into one of following levels.

- **Device level** This is an estimate about the amount of energy the whole mobile device has consumed when used. This is the only level which can be also measured directly so this level of estimate does have an important role when evaluating the accuracy of the model type.
- **Application level** Application level is about estimating the consumption of a single process or application. This is useful for end users who want to evaluate energy efficiency of an application to save battery or to prefer energy efficient applications.
- **Thread level** Thread level is estiamting the consumption of a thread inside an application. This is also very useful for application developers. For example Pathak et al. [8] noticed, when evaluating their program, that Angry Birds, which is very popular commercial game,

does use third party tracking module which is responsible for most of the consumption and is was easy to identify with their tool because the module was running on separate thread.

- **(System) call level** Call level refers to pinpointing the consumption to a single system call or another function. This is most useful for application developers who would like to identify which parts of the program exactly are the bottlenecks.

5.2 Estimate types

The estimation type defines what features of the system does the estimation program monitor and how the consumption is linked to monitored features. One possible classification is utilization based and event based.

5.2.1 Utilization based

The utilization based method is about monitoring the utilization of the components and linking it directly to the consumption. At least PowerTutor and Sesame does estimate the consumption by monitoring the utilization of hardware[12][4].

This is very simple and the intuitive approach since the consumption is directly bound to activity levels of components, for example transmitted bytes in second. However, usefulness of this method is very limited. It is obvious that monitoring only the activity level of component does not reveal what exactly caused the activity, so this can only estimate consumption at device level.

5.2.2 Event based

Event based method could be considered as an extension to utilization based method. The idea to estimate the consumption based on activity of the components is still the same but instead of just monitoring the activity level the changes are detected by monitoring hardware events which causes the change on activity.

This method does have great advantages. The first and the most obvious one is that if we notice an event triggered by an action, for example breakpoint on a system call, it is possible to trace what exactly caused the event. For example AppScope uses KProbes, which are dynamic breakpoints offered by linux kernel for debugging, to trace system calls and monitors Android Binder remote procedure calls to known when the components are used and by who[3]. This allows linking the changes on consumption to exact point on the program. This kind of method can be used to bind consumption to specific application, thread and point on the code where operation was initiated so it is possible to get very fine-grained estimate about the consumption.

The second advantage is that event based method allows capturing asynchronous component power behavior, which is not directly related to utilization of component. For example sending many packets can trigger the component to change into higher power state but when the sending has been finished the component might remain on this state a while to prevent unneccessary transitions between states which could cause delays and more energy consumption[2][8]. The event based method can capture this kind of behavior easily since the system call, which is used to actually change the power state, can be monitored.

However it is worth noting that even if the exact cause of consumption could be tracked down to call made by an application it is not always completely clear which application caused the consumption because of the asynchronous behavior. For example when location is requested the GPS will be turned on and it will remain on for a while in case the location is requested again. The request can be traced easily to a specific location and events to enable and disable GPS interface by an operating system can be traced as well. In this case the consumption is easy to link to one application. However the situation gets more complicated if another application would request the location after the first one. The second request wouldn't cause as much consumption as it would if requested separately, because the interface was already turned on, but it does affect how long the interface will remain on. In this case it is hard to say how the consumption should be calculated.

The AppScope does handle this by logging when application started transmitting data and how much data it transmitted until the component power state changes back to idle. After that it estimates the consumption of component from whole timespan. If multiple applications were using the same component simultaneously it divides the consumption to all applications which were using the component at same time. The division is weighted by amount of data each application transmitted.[3]

5.3 Requirements

The different ways of measuring and estimating the consumption require different levels of access to a device. These are not really characteristics of an estimation model but consequences of certain choices. These are presented here because they have important role in deciding what kind of technique to use. In this paper, software access defines how much control the estimation process demands from the device.

5.3.1 System Requirements

The operating system requirements in this paper refers to intrusion level required by the tool.

- **User space** This is the most user friendly and the most constrained choice. It means that the application can be run on any device without having to root[3] the device. It allows use of the program by any supported device without modifications. However this is usually possible only with utilization based models.

- **Rooted access** Rooted access means access to a full device without restrictions. The user is able to access or modify the system without any limitations which normally exist to prevent malware and to protect digital rights. In this context, this means that there has to be

[2] This kind of consumption is called tail energy.

[3] Get full access to the device without restrictions

access to a whole device to run software but the system does not have to be modified.

- **Modifications** This is like rooted but the difference is that this also requires modifications to the operating system or libraries. This is the most allowing choice but restricts users to those who have full access to the device, enough knowledge and take the risk of breaking the operating system.

5.3.2 Application Requirements

Some of the estimation techniques, like Eprof[8] and PowerProf[6] does require access to an application source code. Although there is also version of Eprof which relied on modified API which allowed estimation application to work without source code on parts which ran on virtual machine[8]. PowerProf does require source code because the way it works is to instrument the application source code with API.

6 Conclusions

There are many papers published about energy consumption estimation on mobile devices. All of them does have a bit different approach to a subject and some of them present the topic quite differently. However these all present solution to a same problem and the general structure is almost the same in all of them. The real differences are usually an improvements to some part of the process.

References

[1] *Mobile Device Power Monitor Manual*. Monsoon Solutions Inc., 2014.

[2] A. Carroll and G. Heiser. An analysis of power consumption in a smartphone. In *USENIX annual technical conference*, pages 1–14, 2010.

[3] Chanmin Yoon, Dongwon Kim, Wonwoo Jung, Chulkoo Kang, Hojung Cha. AppScope: Application Energy Metering Framework for Android Smartphones using Kernel Activity Monitoring. In *the 2012 USENIX Annual Technical Conference (USENIX ATC'12)*, volume 3, June 2012.

[4] Dong Mian, Zhong Lin. Self-constructive high-rate system energy modeling for battery-powered mobile systems. In *MobiSys 11*. ACM, June 2011.

[5] Jung W., Kang C.,Yoon C., Kim D., Cha H. . Non-intrusive and online power analysis for smartphone hardware components. In *MOBED-TR-2012-1*. Yonsei University, 2012.

[6] Kjaergaard Mikkel Baun, Blunck Henrik. Unsupervised power profiling for mobile devices. In *Mobiquitous*, 2011.

[7] Pathak A. , Hu Y.C. ,Zhang M., Bahl P. ,Wang, Y.M. Fine-grained power modeling for smartphones using system call tracing. pages 153–168, 2011.

[8] Pathak Abhinav, Hu Y. Charlie, Zhang Ming. Fine grained energy accounting on smartphones with eprof. volume 10-13. ADM, April 2012.

[9] Wikipedia. Iphone — wikipedia, the free encyclopedia, 2015. [Online; accessed 14-April-2015].

[10] Wikipedia. Lithium-ion battery — wikipedia, the free encyclopedia, 2015. [Online; accessed 14-April-2015].

[11] Wikipedia. System on a chip — wikipedia, the free encyclopedia, 2015. [Online; accessed 14-April-2015].

[12] L. Zhang, B. Tiwana, Z. Qian, Z. Wang, R. Dick, Z. Mao, and L. Yang. Accurate online power estimation and automatic battery behavior based power model generation for smartphones. In *Proceedings of the eighth IEEE/ACM/IFIP international conference on Hardware/software codesign and system synthesis*, pages 105–114. ACM, 2010.

User trajectory recognition in an indoor environment

Sami Karvonen
Student number: 78884T
sami.karvonen@aalto.fi

Abstract

This paper is a detailed survey of state of the art methods for user trajectory recognition using inertial sensors on commodity smartphones. The basic components needed for indoor tracking are an accelerometer, gyroscope, digital compass and the floor plan of the building. User trajectories are useful for example for navigating inside a building and also applications of targeted advertising are possible. Recording user trajectories is separated into three different tasks which are step detection, step length estimation and heading detection. The main challenges in these tasks are small errors in the sensors which can accumulate into significant errors in the user's path. Possible solutions to these errors are particle filters. Other challenges come from user's individual differences in usage patterns of their smartphones and differences in walking patterns.

KEYWORDS: Wi-Fi fingerprinting, indoor localization, trajectory recognition, step recognition, heading recognition, particle filters

1 Introduction

Indoor navigation is a useful technology having many applications such as navigating inside a building or locating people in medical emergencies and many more. Since GPS does not work inside a building, multiple different solutions has been developed for indoor localization. Applications using technologies such as WiFi-fingerprinting[3], GSM-fingerprinting[10] and methods using the Earth's magnetic field[9] have been developed. These methods use prerecorded fingerprint databases for localization, which requires a lot of manually gathered data. The methods also have problems with differently calibrated or inaccurate sensors on mobile devices.

This paper focuses on user trajectory tracking using inertial sensors such as an accelerometer and gyroscope. User trajectory is recorded as a displacement from a known starting point and floor plans are used to assist in determining the user's path inside the building. When the path is known, it can be used to establish the user's current location inside the building; also, gathering fingerprint data along the path is effortless. For example, the Zee system [8] employs user trajectory recognition to crowdsource the gathering of Wi-Fi fingerprint data.

User traces are mainly used for indoor navigation. They can be used to help current Wi-Fi fingerprinting-based applications. These applications require an extensive fingerprint database which is very time consuming to collect. Recording the steps between fingerprints to measure a distance between them can reduce the number of separate fingerprints needed for deploying indoor navigation system[11]. Indoor navigation is useful in itself but it also enables many different kinds of applications which rely on it. For example, shops inside a mall could use the data for a targeted advertising system which sends discount advertisements to people who are just walking past the shop. Other applications could be to organize the evacuations of a building in an emergency situation.

The background section introduces the main components, some reasons why the trajectories are useful and the main challenges for recording user trajectories. The solution section provides solutions for the main challenges such as those encountered with accumulating sensor errors and differences in usage patterns of the mobile device. The solutions also introduces some algorithms from different papers and compares how their solutions differ from each others.

2 Background

This section introduces the key components necessary for user trajectory tracking and the main challenges involved.

2.1 Key components for trajectory tracking

Mobile devices have multiple inertial sensors to detect the phone's movement and orientation. A 3-axis accelerometer detects the acceleration of phone, which is very important for detecting the steps of the user. When we add a compass and gyroscope it, becomes also possible to detect the direction of the movement. With these three sensors user trajectory recognition is possible but in reality the sensors of a typical off-the-self smartphones are inaccurate and the errors in them makes path recognition difficult.

In addition to the phone's sensors, a floor map of the building in which the user is moving is also important for two reasons. Firstly, the map is obviously necessary to show the user his location in the building, and secondly, the map can be used in combination with the sensor readings to rectify any mismatch between the real path and that shown by the sensor readings. For example, if a user is walking along a straight corridor, and the sensor readings indicate a slightly different trajectory, such as walking through a wall, which is impossible, there must be an error. This error can be corrected by comparing the path indicated by the sensors and the position of walls shown by the map.

The actual recording of a user trajectory requires knowledge of two things. First, we need the heading of the user.

For this purpose we need the built-in compass to determine the heading and the gyroscope to determine the phone's orientation. Secondly, we need to determine the distance travelled. In dead reckoning, the distance is determined by the steps taken by the user. Each individual step can be recognized by using the phone's accelerometer, but since steps are not a very reliable distance metric, we also need to estimate length of each step. Now that we know the heading of each step and the length of the steps, we can calculate the full trajectory of the user.

Wi-Fi fingerprints are used for a different kind of indoor navigation where the Wi-Fi fingerprints are first manually recorded inside the building and then used to locate the user. The goal of user trajectory recognition is to render it unnecessary to gather this fingerprint database but we can still make use of the technique. We do not have this database to start with, but at any time when tracking the user's path inside the building, we can record the fingerprint data and use it in the future as a known starting point in the building in order to make the trajectory tracking more precise. The fingerprints provide points in the building where the user's position is know, and these points can be used to guide the tracking of the user's trajectory. Also, they might be helpful in correcting errors in step length estimation by providing two known points with a known distance between them. This would require thought, that the we could use the map to measure the distance between the points, because it is impossible to calculate the distance between two fingerprints by the fingerprint data alone.

2.2 Challenges

User trajectory recognition usually starts with detecting the user's steps and their direction. Since mobile devices such as phones and tablets are used in very different manners, it is some times hard to detect delicate movements such as steps. Some people browse the web while they walk and others just keep it in a pocket or a bag. All these different locations have to be taken into account while detecting the user's steps, moreover, random movements such as putting the phone in the pocket can be hard to differentiate from steps. Also steps are not very accurate measure of length, thus algorithms to measure the length of each step in different situations are necessary.[5]

The inertial sensors which mentioned above are not very accurate and even a small error might be a problem. A small error in a sensor is not a big problem in itself but when we are continuously tracking the user's trajectory, even a small error can accumulate over the iterations and become very a significant error in the overall path of the user [8]. For this reason we need some way to refine the sensor readings.

The digital compass is necessary for detecting the users heading. In typical off-the-self mobile devices it is even more inaccurate than the inertial sensors. Also a lot of magnetic interference affect it, which makes it hard to make use of in a useful manner. In most cases this is the most difficult challenge to solve in trajectory recognition.

3 Solutions

User trajectory recognition is a difficult challenge in an indoor environment and it is easier to solve it when it is divided into smaller independent tasks. This section surveys the state of the art solutions to these tasks and compares them. Usually User trajectory tracking is divided into four tasks: step detection, step length estimation, heading detection and filtering out the errors from sensors.

3.1 Step Detection

The basis for step detection on a smartphone comes from the embedded accelerometer sensor. The readings from the sensor are recorded to a buffer and generated into a graph. Figure 1 shows these similar patterns, which can be recognized as steps. The problem with reliably recognizing steps on the graph is that the phone is also subject to other motions such as the user sitting down or turning. These motions can be similar to steps in the accelerometer graph and may lead to false positives on step detection. [6]

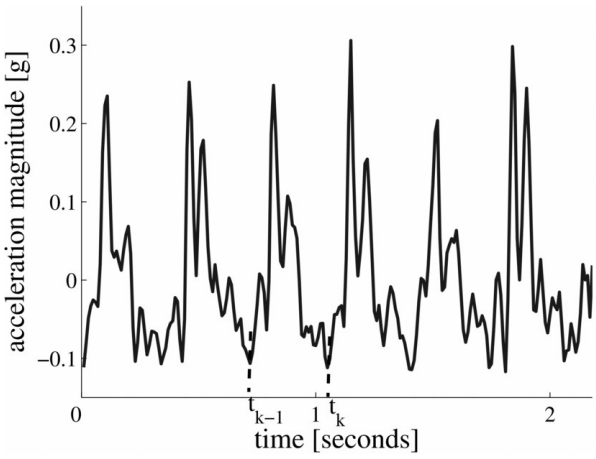

Figure 1: Step graph from accelerometer. Source [6]

Detecting steps reliably is an important part of tracking user trajectory and many different solutions have been proposed. Y.Liu et al. [6] propose a solution which compares neighbouring sequences in the graph and if they are similar enough they are recognized as steps. A slightly different solution in [5] which focuses on detecting the peaks and valleys of the graph and filtering out false peaks using a given threshold. Also some heuristics are added to decrease the number of false positives. In my opinion, these solutions are not mutually exclusive and might actually benefit from each other if used together.

In trajectory recognition, steps are also used to measure the distance which the user has travelled but the problem is that step length can vary greatly in different circumstances. Obviously, different people have different step lengths, but it can also differ depending on a person's health status or the type of the ground the person is walking on, and on many other different factors [5]. Some generalized algorithms for step length estimation have been proposed in [4] but these

algorithms require too much information, such as user foot length, and make too many assumptions such as assuming the phone is located in the user's pocket. These assumptions make the algorithms too inconvenient for use in an actual indoor localization application. This is why we need a more suitable way to estimate step length. One approach is to use particle filters to gather data from the user walking patterns and continuously adapt to changes in step length. Particle filters are discussed more thoroughly in section 3.3.

A.Pratama et al. [7] present an experimental setup to evaluate four equations(1),(2),(3),(4) for step length estimation. The experiment was performed over the distances of 10, 20 and 30 meters by 15 test subjects. The static method(1) assumes that step length can be calculated from the users height.

$$step\_size = height \cdot k \tag{1}$$

The Weinberg(2) approach assumes that the vertical bounce of walking is proportional to the length of the steps.

$$step\_size = k \cdot \sqrt[4]{a_{\max} - a_{\min}} \tag{2}$$

The Scarlet approach(3) considers that the length of the steps can be calculated from the spring of the steps.

$$step\_size = k \cdot \frac{\frac{\sum_{k=1}^{N} |a_k|}{N} - a_{\min}}{a_{\max} - a_{\min}} \tag{3}$$

The Kim approach proposes an equation which represents a relation between step length and average acceleration during the step.

$$step\_size = k \cdot \sqrt[3]{\frac{\sum_{k=1}^{N} |a_k|}{N}} \tag{4}$$

As we see in table 1, the Scarlet method produces the best results in this experiment and the average error is only 1.3913 meters. This level of accuracy should be enough for at least some applications of indoor localization. This experiment was performed with the assumption that the user is holding the phone in his hand. This is not always the case in the real world, thus it would be interesting to see how these equations perform when the phone's location is not known.

Table 1: Results from the experiment in paper [7]

Method	Estimation error average (m)
Static	2.5815
Kim	1.6917
Weinberg	1.4357
Scarlet	1.3913

3.2 Heading Detection

Determining the heading of each step is as important for dead reckoning as detecting the steps and their length. Smartphones have a built-in compass which can be used to determine the heading of the phone although it is even more prone to errors than the accelerometer. These errors are caused by objects that effect the surrounding magnetic fields, such as power lines and iron reinforcements in buildings. According to [2], the typical accuracy of the built-in compass is 5 degrees or less. At the level of single steps this is not much, but these errors can accumulate over multiple steps and cause significant errors in the position after multiple iterations.

D.Gusenbauer et al. [2] attempts to solve this problem by presenting an adaptive filtering function. The function relies on the fact that the faster the user is moving, the less likely he is to change direction. Thus, false changes in direction can be filtered out when the user is moving fast, and the filtering can be tuned down when the movement is slower. However, this does not help the fact that at turning points of the movement, such as intersecting corridors, we need to change direction and the errors in the sensor might lead to an error in the estimation of the angle of the turn.

A different approach to this problem is introduced in [5]: it uses particle filters and floor plans to correct the errors in the sensors. More about particle filters in section 3.3.

3.3 Particle filtering

Particle filtering is a general way to filter out error in sensor readings. Simply, in a user trajectory context, particle filtering uses random particles to represent the user's position. The particles move according to data acquired from the sensors. Each of these samples moves independently and the information obtained from a floor plan of the building is used to weight the samples. For example, when a particle hits a wall, it will die and an another particle is generated near a living particle. When this process goes through multiple iterations, the particles that collide with walls because of sensor errors die out, and new particles are spawned near living particles which represent a possible path of the user. The combined weight of the living particles represents the probable position of the user and the impossible paths are filter out when particles die. Using a particle filter can significantly improve user trajectory tracking compared to using only available sensor data. Figure 2 shows the improvement made to an estimated path when particle filtering is applied to the sensor readings. [1]

Particle filters can be also used in a different way to improve user trajectory recognition. An approach to using particle filtering for step length personalization is proposed by F.Li et al. in [5]. The basic idea is the same as before, but in this case a personalization model is applied to each particle as it is generated. This works well in long corridors where the main reason for particles to die is the fact that the personalization model makes the particles collide with a wall– in other words, when the particles do not match the real path they die out. This makes the particles of the less accurate models disappear, which, in turn, means the aggregation of the remaining particles represent a more accurate model of reality. A problem arises at turn points where particles die because of the turn and not the model. This is why it is important to also implement a turn detection algorithm and use it to decide if a particle died because of a turn or its personalization model.

Particle filtering is a powerful method which have been used in many applications of user trajectory recognition. For

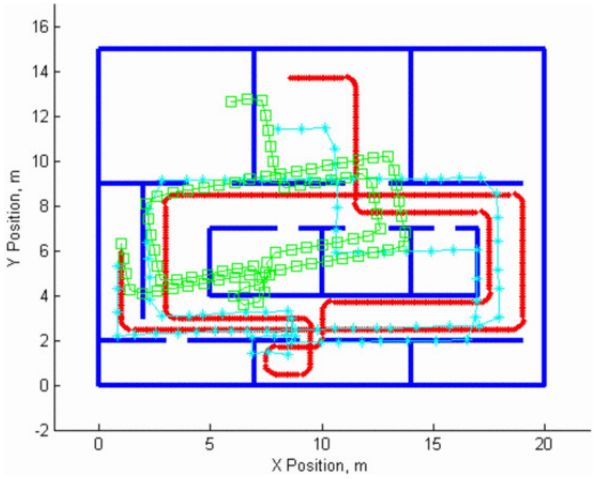

Figure 2: Walking simulation. Blue lines are building walls. Red(continuous line) is the real path. Green(squares) is the path with only sensor data. Cyan(stars) is estimated path when particle filtering is applied. Source [1].

example, F.Li et al. [5] use it for step length personalization and P.Davidson et al. in [1] use it for overall error correction in the path. The problem with particle filtering is that it is a simulation method which requires a lot of processing power. This is because every particle's path needs to be separately calculated. Since we are working in a mobile environment the power requirements should be taken into account and careful consideration of where the high power requirements are necessary and where they are not.

3.4 Observations of different solutions

The step detection part of the problem seems to be sufficiently solved in many of the research papers. For example, F.Li et al. [5] report their step detection algorithm to have only 1.6% of error and A.Rai et al.[8] report on 0.6% error. The comparison of step length estimation algorithms presented before shows that a 30 meter distance can be measured by steps with an average error of 1.4 meters. The overall error in the path length is a combination of these. The combined error margins should be small enough for a working application of indoor trajectory recognition. Thus this part of the problem should be considered solved.

The problem comes with accurate heading detection. In many occasions [8] [5] [1] it is mentioned that the error of the digital compass in typical off-the-self mobile devices is very high. It can be as high as 12 degrees in average. This is the main source of errors in the path in most of the applications of trajectory recognition. P.Davidson et al. in [1] seems to be successful in correcting these error with particle filtering but the data in their research paper is not very extensive, thus further research is needed.

4 Conclusion

This paper surveys the state of the art indoor trajectory recognition techniques and presents the main challenges on recording user traces. User trajectory tracking in an indoor environment is a challenging problem with many separate tasks that all affect the accuracy of the whole system. The three main tasks are step detection, step length estimation and heading detection which all require data from the fairly inaccurate sensors of a smartphone. There is a lot of research on the subject offering many solutions to these problems but many of the systems still seem inaccurate as a whole. The most challenging problem seems to be the heading detection because of the in accuracy of the digital compass and large amount of magnetic interference which also affect the compass readings. Errors in compass readings are challenging because they accumulate over time and cause a large drift in the overall path. For future work more research about the heading detection compass accuracy is necessary. The main questions that require research are: is there any way to refine the data from the digital compass to make it more accurate, and does particle filtering work well enough in every situation to correct the errors of inaccurate compass readings? Further more, some larger scale practical experimentation of the solutions to each separate problem would be useful to determine which of them work best, and to combine the efforts made in all of them.

References

[1] P. Davidson, J. Collin, and J. Takala. Application of particle filters for indoor positioning using floor plans. In *Ubiquitous Positioning Indoor Navigation and Location Based Service (UPINLBS)*, pages 1–4, October 2010.

[2] D. Gusenbauer, C. Isert, and J. KroÌĹsche. Self-contained indoor positioning on off-the-shelf mobile devices. In *Indoor Positioning and Indoor Navigation (IPIN), 2010 International Conference on*, pages 1–9, Sept 2010.

[3] M. N. Husen and S. Lee. Indoor human localization with orientation using wifi fingerprinting. In *Proceedings of the 8th International Conference on Ubiquitous Information Management and Communication*, ICUIMC '14, pages 109:1–109:6, New York, NY, USA, 2014. ACM.

[4] J. Jahn, U. Batzer, J. Seitz, L. Patino-Studencka, and J. GutieÌĄrrez Boronat. Comparison and evaluation of acceleration based step length estimators for handheld devices. In *Indoor Positioning and Indoor Navigation (IPIN), 2010 International Conference on*, pages 1–6, Sept 2010.

[5] F. Li, C. Zhao, G. Ding, J. Gong, C. Liu, , and F. Zhao. A reliable and accurate indoor localization method using phone inertial sensors. In *the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*, pages 421–430, September 2012.

[6] Y. Liu, M. Dashti, and J. Zhang. Indoor localization on mobile phone platforms using embedded inertial sensors. In *Positioning Navigation and Communication (WPNC), 2013 10th Workshop on*, pages 1–5, March 2013.

[7] A. Pratama, Widyawan, and R. Hidayat. Smartphone-based pedestrian dead reckoning as an indoor positioning system. In *System Engineering and Technology (ICSET), 2012 International Conference on*, pages 1–6, Sept 2012.

[8] A. Rai, K. K. Chintalapudi, V. N. Padmanabhan, and R. Sen. Zee: zero-effort crowdsourcing for indoor localization. In *the 18th annual international conference on Mobile computing and networking (Mobicom '12)*, pages 293–304, August 2012.

[9] K. P. Subbu, B. Gozick, and R. Dantu. Locateme: Magnetic-fields-based indoor localization using smartphones. *ACM Trans. Intell. Syst. Technol.*, 4(4):73:1–73:27, Oct. 2013.

[10] A. Varshavsky, A. LaMarca, J. Hightower, and E. de Lara. The skyloc floor localization system. In *Proceedings of the Fifth IEEE International Conference on Pervasive Computing and Communications*, PERCOM '07, pages 125–134, Washington, DC, USA, 2007. IEEE Computer Society.

[11] X. Zhang, Z. Yang, C. Wu, W. Sun, Y. Liu, and K. Liu. Robust trajectory estimation for crowdsourcing-based mobile applications. *Parallel and Distributed Systems, IEEE Transactions on*, 25(7):1876–1885, July 2014.

Virtual Machine Consolidation with Multi-Resource Usage Prediction

Kimmerlin Maël
Student number: 399724
mael.kimmerlin@aalto.fi

Abstract

One of the major concerns currently in cloud computing is the energy efficiency in the use of the data centers. In this context, cloud providers aim at avoiding underutilization and overutilization of their infrastructure thus reducing both the power consumption and the Service Level Agreements violations(SLAs). The live Virtual Machines migration technology enables the consolidation of virtual machines, thus allowing cloud providers to reallocate virtual machines into a few servers and switch off unused physical machines. depending on the load level in the current host. However, the migrations have to be triggered carefully due to the diversity and the actual variability of the virtual machines workload. This paper focuses is on the resource utilization prediction and on the overload or underload detection. We survey several efficient paradigms that are effective in forecasting the next values, such as Linear Regression and Neural Networks, to estimate the prediction function based on the history of the host. We then present a virtual machine consolidation algorithm with prediction for improving the energy efficiency of cloud data centers. We use the CloudSim simulation toolkit to evaluate our proposed solution.

KEYWORDS: VM consolidation, load prediction, CloudSim, energy efficiency

1 Introduction

Data centers and cloud computing raise several concerns, mostly about energy efficiency. A problem associated with the high speed development of Information and Communication Technologies (ICT) is the high carbon footprint it generates. Indeed, this footprint forms a relatively high part of the global carbon emissions and concerns have been raised about the efficiency of data centers that are always overprovisioned to sustain peak demand [12]. Furthermore, overprovisioning a data center effects the Quality Of Service (QoS), i.e. the response time, of the hosted applications. Therefore, cloud providers need to fulfill the Service Level Agreement (SLA) for assessing the QoS level offered to their customers. On the other hand, underprovisioning consumes resources uselessly resulting in unnecessary activation of physical servers, thus increasing the actual costs.

The main challenge is to decide whether a host is overloaded or underutilized. When a host is underutilized, all virtual machines (VMs) from this host are selected for migration to suitable hosts in order to turn off the idle hosts to save energy. [3]. Specifically, underutilized host migrations should not impact the current fully utilized hosts nor should they create overloads on the other hosts. Similarly, when a host is overloaded, the orchestrator should migrate some VMs to reduce the load [6]. The orchestrator can trigger both underutilized and overutilized migrations by using threshold values so that when the resource consumption reaches a low or high threshold, it starts migrating VMs [1]. However, the current CPU or memory load is not an adequate indicator as it might be biased by temporary peaks in the VM resources consumption. Thus it may cause unnecessary migrations leading to an increased load on the hosts and network. Consequently, it is important to base the migration decision on the trend and take into account the history, rather than short events. The orchestrator can predict the resources utilization in order to trigger migrations only for hosts that are effectively overloaded or underloaded. Moreover, Beloglazov et al. have shown that maximizing intervals between relocations improves the VM consolidation [3]. However, even if resource prediction has been emphasized lately in research, it mostly takes into account a single CPU resource. Mastroianni et al. have shown that it is more efficient to take several resources such as CPU, memory and bandwidth utilizations, into account since in cloud data centers CPU is not the only critical resource. [12] Some hosts (such as caches or in-memory databases for example) store the most requested content in memory and serve them on demand, thus the requirements for such hosts are on the memory and not on the CPU, as it does not require intensive computations. Other hosts could have their main requirement ont he bandwidth, such as web servers for which the bandwidth is as important as the CPU capacities. Thus basing the migration decision on CPU only does not include all the parameters. A host might not be over loaded considering the CPU, but its memory could be depleted, thus creating SLAs violations because of the swapping or because of the killing of the processes. Thus it is necessary to take all the parameters into account.

The main contributions of this article are as follow.

- We survey the current research on prediction, based on history for multiple resources, and on overloaded and underloaded server detection.
- We then present an efficient multi-resource usage prediction (MUP) approach to estimate the short-term and long-term future utilization in terms of multiple resource types based on the history of those resources.
- Using MUP, we propose an efficient multi-resource overloaded and underloaded host detection mechanism based on history.

- We perform simulations to validate the proposed solution.

The remainder of the article is organized as follows. Section 2 discusses the state-of-the-art research performed on resource utilization prediction and on overload and underload detection. Section 3 details a study of the methods to predict the next utilization level of the resources of the servers, such as linear regression and neural networks. Section 4 presents the detection mechanisms of overloaded and underloaded hosts for efficient VM migration. Section 5 presents the experimental setup and compares the results obtained by our prosed solution with existing solutions for VM consolidation. Finally, Section 6 concludes the article.

2 Related Work

Several approaches have been taken to solve the problem of overloaded or underloaded host detection. One of the simplest approaches was proposed by Gmach et al. [7]on VM consolidation. In detail, the authors set a static threshold and used the current CPU utilization to determine whether a host is overloaded. Consquently, the static threshold is set to 85% based on statistical analysis of the workload traces. However, setting static threshold and using the current resource usage are not efficient in environments with dynamic workloads, as the short-terms variations are considered and cause unnecessary migrations. Beloglazov et al. proposed an algorithm to dynamically adapt the value of the threshold [2]. In order to make more accurate predictions of the workload, the authors proposed a model based on Markov Chains that enables them to formalize the placement problem using QoS metrics [3]. However, this approach is not adequate for variable workloads, thus they adapted the work of Luiz et al. on the Multisize Sliding Windows to predict the next workload more accurately [11]. This approach is limited to a single resource. It does not consider other parameters that might cause and overload on servers. Thus, Mastroianni et al. proposed another approach combining multi-resource placement problem using ant algorithms to solve the complex bin packing problem with multiple resources using multiple simple operators [12]. This approach is rather complex, as the bin packing problem is at least NP-hard and requires high computation power. Thus we consider simpler approach by using load prediction to temper the current utilization statistics.

3 Predicting Resource Utilization Based on History

In statistical analysis, several methods exist to predict the next value of a set of values such as ordinary least squares, logistic regression, and vector auto-regression models. This section presents two naturally efficient and effective forecasting paradigms, namely linear regression and feedforward neural network, to estimate the prediction function according to the history of the VM past resource utilization. In addition, we incorporate sliding window technique in the training and long-term prediction strategy.

3.1 Linear Regression

Linear regression consists of modeling the relationship between one or more input variables and the output variable. For example, the CPU load can be predicted using a simple linear regression (SLR) by taking the CPU load of time t as input variable and the CPU load of time $t+1$ as output. SLR can be written as in Eq. (1).

$$y(x) = \alpha + \beta x + \epsilon_x \tag{1}$$

ϵ_x represents the error for load x, equivalent to the difference between the predicted and the exact values; α and β are the regression coefficient parameters that are estimated according to the resource history in order to minimize the overall error ϵ, using the least square method.

The original least squares (OLS) method is the most popular method to estimate α and β parameters. The OLS is defined with a data set of $N+1$ values by minimizing the sum of the square errors as in Equation 2.

$$S(\alpha, \beta) = \sum_{i=0}^{N} (\epsilon_{x_i})^2 \tag{2}$$

$$S(\alpha, \beta) = \sum_{i=0}^{N} (y(x_i) - \alpha - \beta x_i)^2 \tag{3}$$

The two gradients of the function should be equal to zero as shown in Equations 4 and 5.

$$\frac{\partial}{\partial \alpha} \sum_{i=0}^{N} (y(x_i) - \alpha - \beta x_i)^2 = -2 \sum_{i=0}^{N} (y(x_i) - \alpha - \beta x_i) = 0 \tag{4}$$

$$\frac{\partial}{\partial \beta} \sum_{i=0}^{N} (y(x_i) - \alpha - \beta x_i)^2 = -2 x_i \sum_{i=0}^{N} (y(x_i) - \alpha - \beta x_i) = 0 \tag{5}$$

Thus α and β can be defined as in Eq. (6) and (7) as follows.

$$\alpha = \bar{Y} - \beta \bar{X} \tag{6}$$

$$\beta = \frac{\sum_{i=0}^{N} (x_i - \bar{X})(y(x_i) - \bar{Y})}{\left(\sum_{i=0}^{N} (y(x_i) - \bar{Y}) \right)^2} \tag{7}$$

where $\bar{X}$ and $\bar{Y}$ are the average values of the series of x and $y(x)$. This method has been extensively described by Farahnakian et al.[5]

In our prediction model, we focus on multi-resources, i.e. CPU, memory, network, load prediction, thus we consider a Multiple Linear Regression (MLR). The principle of MLR method is similar to SLR but we have more than one input variable. In the MLR prediction function, all input can be combined to define an overall load of the host. Then MLR could be used to predict the future overal resource load of the host based on the current load of all resource types. However, the relationship between multiple types of resources is not linear due to the diverse set of users applica-

tions, i.e. general-purpose applications include a balanced amount of CPU, memory, storage and network resources; computer-intensive applications need more CPU resources; and database and memory caching applications need more memory resources. So we consider multiple simple linear regressions. Therefore, CPU, memory, storage or network load, can be computed separately. Then the overload or underload would be determined considering each indicator separately, with separated triggers.

3.2 Feedforward Neural Network

Neural networks allow non-linear predictions and are designed as a brain, with neurons and synapses. The neurons are entities able to perform simple operations based on simple conditions. In our prediction model, we use the full-connected three layer feedforward neural network (FFNN) [8]. The number of input neurons correspond to the input data, the number of output neurons to the output data and the number of hidden neurons can be selected empirically.

The neurons of FFNN have several characteristics, the weight of the links connecting them and their activation function. An error-correction methods enable to tune the weights. The activation function computes the output of the neuron considering the inputs. Usually the activation function is a sigmoid that outputs a value between 0 and 1. The sigmoid activation function is defined by Rumelhart et al. in [15] as :

$$f : \mathbb{R} \to \mathbb{R}, f(x) = \frac{1}{1 + e^{-x}} \tag{8}$$

Let $w_{i,j}$ the weight of the link between neuron i and neuron j, i.e. w_{I_3,H_2} is the weight of the link between neuron 3 of the input layer and neuron 2 of the hidden layer. There are several learning methods for neural networks. Some of them takes advantage of non-linearity, some samples the data sets randomly. During the training, we use the back-propagation approach to estimate the synaptic weights of the network that minimize the sum of squared errors in the output neuron. The iterations in the training phase are discussed in the following. [14]

First, the weights are initialized randomly. Then for every set of data in the learning phase $x_j(i)$, i being the identifier of the data set and up to N, j the identifier of the input neuron, the network is fed with each set of data and an output value is calculated for all the neurons. Then the output of each neuron is computed using the Eq. (9).

$$o_n(i) = f\left(\sum_{j=1}^{H} w_{n,j} x_j(i)\right) \tag{9}$$

We then calculate $E_n(i)$ as the difference between the real output $\bar{O_n}$ and the calculated output, and we update the weights with the formula given in Eq. (10).

$$w_n(t+1) = w_n(t) - \gamma \sum_{i=1}^{N} E_n(i) * f'\left(\sum_{j=1}^{H} w_{n,j} x_j(i)\right) \tag{10}$$

The condition to halt the training of a neural network is usually based on the sum of the squared errors that should be under a defined limit. γ represents the learning rate of the network. This value should be chosen carefully as if it is too low the network will converge too slowly and if it is too high, the network might oscillate around its optimal values. Once the network was adjusted with all the input data set of the training, it can be given current data set to predict the next ones.[10].

3.2.1 Sliding Window Prediction

The sliding window technique maps multiple input sequence to an individual output by using a prediction model. For example, when predicting the value for time $t + 1$, the values used are from the history from $[t - k; t]$ for a sliding window of interval k [10]. Thus the size of the problem remains reasonable, even if the size of the history increases.

3.3 Long-term Prediction Strategies

In order to make consistent decisions about the load induced by a VM, a long-term prediction mechanism is needed. There are several strategies to extend the short-term prediction methods for long-term predictions, as described in [8]. We consider that the prediction function as presented in Eq. (11)

$$U_{x,t+1} = f\left(U_{x,t}, U_{x,t-1}, ...\right) \tag{11}$$

3.3.1 Recursive Strategy

Recursive strategy predicts n steps ahead separately, each of the predictions would include all the previous predicted values as described in Eq. (12). This method may increase the errors through the prediction steps.

$$U_{x,t+n+1} = f\left(U_{x,t+n}, U_{x,t+n-1}, ..., U_{x,t}, U_{x,t-1}, ...\right) \tag{12}$$

where t is the current time, U the resource utilization of resource x and n the number of steps ahead for the prediction. All the resource utilization level in the future are calculated using the formula. Thus it is a recursive approach.

3.3.2 Direct Strategy

Instead of predicting n steps and using them to predict the $n + 1$, the direct straategy uses only a single step but based on predictors separated by n steps, as presented in Eq. (13). Therefore, the interval between each value of the history grows to match the step between the last real value and the first predicted one. Nonetheless, this model has also important drawbacks, as it requires a longer history and may hide spikes. It is highly probable that short-term spikes occurs during the interval between two samplings and thus are not reflected in the prediction. Thus, this method can provide a good average prediction only if it is fed with enough data. The complexity of this method also increases linearly with

the number of predictions as the history data set has to be changed for every new prediction.

$$U_{x,t+n} = f\left(U_{x,t}, U_{x,t-n}, U_{x,t-2*n}, ...\right) \quad (13)$$

3.3.3 Combined Strategy

The combined strategy, namely DirRect, has been proposed to overcome both issues mentioned above. In detail, the DirRec method uses a direct strategy to predict the first value and a recursive strategy for the next ones as presented in Eq. (14), thus reducing the complexity and the possible error.

$$U_{x,t+k*n} = f\left(U_{x,t+k*n-n}, U_{x,t+(k-1)*n}, ..., U_{x,t}, U_{x,t-n}, ...\right) \quad (14)$$

where k is the number of n steps ahead of the approach.

4 Overloaded and Underloaded host Detection

Predicting the load of a host can be achieved by focusing on multiple types of resources. It is possible to focus on the host load history, for example, and predict the next values out of it. However, this is not reliable in the case of virtualization and VM migration as the load is not continuous. When a VM is migrated to or from a host, the load of the host can increase or decrease greatly. Thus, focusing on the host itself is not the right solution. However, the history of a VM does not change wherever it is placed, thus it is possible to predict more accurately the load of a VM (as long as there are no disruption caused by the migration). Thus it is possible to predict the load of a host by predicting the load of all its VMs. Indeed, the load on a host is equal to the load of all its VMS plus an overhead for the virtualization system and the operating system of the host. Thus, we focus the predictions on the VMs, while focusing the triggers of underload or overload detection on the physical host. We consider the overhead to be not linear with regards to the number and load of VMs. Considering the memory for example, the overhead consists of the memory of the operating system, the memory used by the virtualization system (both fixed amount and variable amount considering the number of VMs and their capacities). In order to determine the order of the overhead, Tong et al. ran experiments to measure the overhead.[16] They found out that the overhead is related to the number of VMs running and has a minimum due to the footprint of the system.

We consider the overloaded host detection and uderloaded host detection with similar approaches. For instance, if a method use a maximum threshold to determine if a host is overutilized by reaching this threshold, a similar approach can be taken for underloaded host detection, using a minimum threshold to determine if a host is underutilized by reaching this minimum threshold. The detection method for overloaded host can be applied for underutilized host. thus we mainly focus on overloaded host detection.

4.1 Parameters correlation

The correlation of parameters such as the CPU, memory, storage or network loads is highly variable depending on the purpose of the VM [12]. For some tasks, VMs require mainly memory (I.e. Hyrise In-memory databases[9]) without needing high CPU performances. Other types require computing power, or storage or bandwidth. However, in particular cases, some parameters can be highly correlated. Thus, we expect that predicting the next steps using all the parameters increases the accuracy of the prediction. Using FFNN would enable us to learn the correlation between the parameters during the training phase. However due to the linearity in the linear regression method, we expect the accuracy of the prediction to decrease when taking the parameters.

4.2 Static threshold

A first method to detect the overload is to use static thresholds as described in [3]. In such cases, a static threshold was used and set to 85to determine whether a host is overloaded. This method can be extended to multiple resources by setting a static threshold for each parameter. However, this method presents some drawbacks, for example, in case of a spike in the load, migrations will be triggered while the load decreases, thus creating more load than needed. Thus improvements of this method are needed.

4.3 Adaptive threshold

The previous method presents a drawback. If the load of the VMs varies widely, it is highly probable that some unexpected peaks happen. Those peaks could create high temporary overload, that was not expected as they happen between the verifications. A solution would be to adapt the threshold. If the VMs are known to have a highly variable load, the threshold can be reduced for the overloaded host detection. Thus in case of a peak, its impact is lower as there are more free resources. Similarly, for underloaded host detection, if the VMs have highly variable loads, the threshold is also increased, so that a temporary reduction of the load does note cause a consolidation that would need to be undone when the load comes back to normal level. Thus the variance of the load is a good indicator to adapt the threshold and perform more efficient migrations, by removing the unneeded migrations. However the energy consumption increases slightly.[13]

5 Performances Evaluation

5.1 Experimental Setup

CloudSim-toolkit is a powerful simulation tool that allows to test seamlessly the performances of a new application or feature as it models and simulates large clouds (from cloud resources to datacenters)[4]. CloudSim-toolkit integrates traces of workload from PlanetLab servers. We used the traces of 10 different days. The results presented here are averaged.

CloudSim-toolkit uses Allocation Policies to determine if a host is overloaded. We extended the static threshold policy to integrate multiple resources and load prediction. We implemented the FFNN and SLR prediction methods for the load prediction with the sliding window method.

5.2 Evaluation Metrics and Comparison Benchmarks

We considered several metrics for comparing our solution to the existing ones. The first metric we used is the energy consumption of the datacenter. The energy consumption varies as some servers are turned off or on to sustain the load. Our aim is to minimize this parameter to reduce the footprint of the datacenters. The second metric used is the number of VM migrations. Migrating VMs consumes resources, thus reducing this numbers helps to reduce the energy consumption. The third metric is the overall SLAs violations. This describes the service provided to the user. The aim is to make it null. Between those three metrics, there is a tradeoff to balance. It is not possible to minimize all of them as when you minimize the SLAs violations for example, the energy consumption increases as more hosts are turned on to avoid overutilization.

We ran several test cases. For the SLR, we used different size of sliding window, from 5 samples to 100 samples. We expect that increasing the samples changes the prediction to values closer to the average load of the VM, thus hiding peaks, and increasing the SLA violations.We also changed the number of predicted steps to study its impact. For the FFNN, we also varied the size of the sliding window from 5 to 100 samples and the number of steps ahead in the prediction. We used two different ways for the prediction using the FFNN. The first one was to consider every parameter separately and use FFNN for each of them separately, thus not capturing the interaction between the parameters. Then we used a single FFNN with all the parameters as input and several output, in order to capture those interactions.

5.3 Experimental Results

5.3.1 FFNN Results

Unfortunately, none of our multiple implementations and test case using FFNN succeeded. In most of the cases, the FFNNs were never converging, even when using high error rates. However, when they were converging, the time needed to train them was far too high considering the time interval between the sampling of the resources utilization. The prediction mechanism was not fast enough to produce usable output. Thus using FFNN in this context was a failure. However, more research is needed before discarding completely this solution.

5.3.2 Simple Linear Prediction Results

The energy consumption is depicted in Fig. (1). When using SLR, the hosts are detected as less overloaded as it hides the peaks in the resource consumption. Thus the energy consumption is highly limited as less hosts are switched on. Similarly, the number of VM migrations is much lower when using the prediction algorithm as shown in Fig. (2).

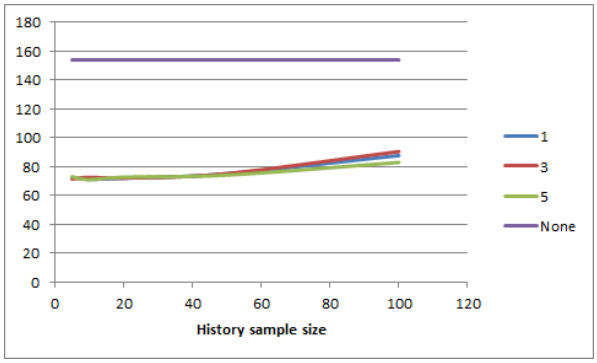

Figure 1: Datacenter energy consumption (kWh)

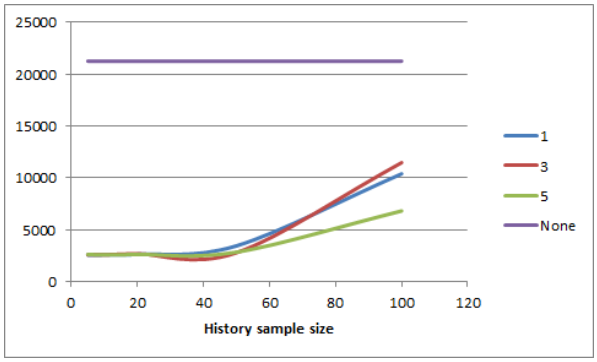

Figure 2: Number of VM migrations

However, the drawback of this high energy efficiency is that the SLA violations percentage nearly doubles as depicted in Fig. (3). Thus this solution is not optimal as the users perceive the resources exhaustion.

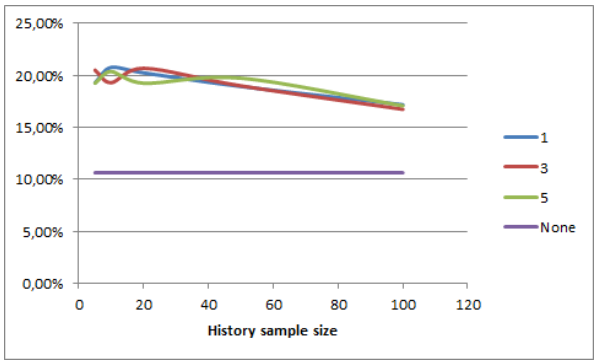

Figure 3: Average SLA violation percentage

5.4 Limitations of the simulation

This simulation presents several limitations. The first main one is that in the traces used from PlanetLab, the load is mainly on the processing capacities, thus the other parameters impact in the overloaded or underloaded host detection is not visible. In order to validate the model based on the CPU, memory and bandwidth, other traces would be needed. However, despite active search for such traces, the PlanetLab traces were the only one available. thus the simulation has a severe drawback on that point.

A second limitation is that the VM selection for the migration process was not taken into account. that is because it was defined as such in the subject, but for sake of completeness it should be included in the study. However, since the same one was used for all the test cases, we believe that it does not invalidate the results obtained.

6 Conclusion

Energy efficiency is the new target in datacenters. Several methods can be used to tend to this objective. The main one is to migrate the VM from underutilized hosts to others while ensuring that none of them get overutilized in order to avoid SLA violations. However, the current method focus on the processing capacities and do not consider other parameters of the physical host. Thus we propose some modifications to integrate the other parameters in order to achieve better consolidation of the VMs. We also present prediction methods to optimize the overloaded and underutilized hosts detection. Those methods are Feedforward Neural Networks and Simple Linear Regression. We implemented both of them and tried them in a simulation using real-world traces from PlanetLab. Even if the FFNN method didn't succeed, SLR gives interesting results in terms of energy efficiency. However, it increases the SLA violations. Thus, further research might be needed to improve further the SLR method by tuning finely its parameters or research with other kind of neural networks.

References

[1] A. Beloglazov and R. Buyya. Adaptive threshold-based approach for energy-efficient consolidation of virtual machines in cloud data centers. In *Proceedings of the 8th International Workshop on Middleware for Grids, Clouds and e-Science*, MGC '10, pages 4:1–4:6, New York, NY, USA, 2010. ACM.

[2] A. Beloglazov and R. Buyya. Optimal online deterministic algorithms and adaptive heuristics for energy and performance efficient dynamic consolidation of virtual machines in cloud data centers. *Concurrency Computat.: Pract. Exper.*, 24:1397–1420, 2012.

[3] A. Beloglazov and R. Buyya. Managing overloaded hosts for dynamic consolidation of virtual machines in cloud data centers under quality of service constraints. *Parallel and Distributed Systems, IEEE Transactions on*, 24(7):1366–1379, July 2013.

[4] R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. F. De Rose, and R. Buyya. Cloudsim: A toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms. *Softw. Pract. Exper.*, 41(1):23–50, Jan. 2011.

[5] F. Farahnakian, P. Liljeberg, and J. Plosila. Lircup: Linear regression based cpu usage prediction algorithm for live migration of virtual machines in data centers. In *Software Engineering and Advanced Applications (SEAA), 2013 39th EUROMICRO Conference on*, pages 357–364, Sept 2013.

[6] H. Ferdaus and M. Murshed. *Energy-aware Virtual Machine Consolidation in IaaS Cloud Computing*, chapter 55, page 33. Springer, 2014.

[7] D. Gmach, J. Rolia, L. Cherkasova, and A. Kemper. Resource pool management: Reactive versus proactive or let's be friends. *Comput. Netw.*, 53(17):2905–2922, Dec. 2009.

[8] A. Grigorievskiy, Y. Miche, A.-M. Ventela, E. Severin, and A. Lendasse. Long-term time series prediction using op-elm. *Neural Networks*, 51(0):50 – 56, 2014.

[9] M. Grund, J. Krüger, H. Plattner, A. Zeier, P. Cudre-Mauroux, and S. Madden. Hyrise: A main memory hybrid storage engine. *Proc. VLDB Endow.*, 4(2):105–116, Nov. 2010.

[10] S. Islam, J. Keung, K. Lee, and A. Liu. Empirical prediction models for adaptive resource provisioning in the cloud. *Future Gener. Comput. Syst.*, 28(1):155–162, Jan. 2012.

[11] S. Luiz, A. Perkusich, and A. Lima. Multisize sliding window in workload estimation for dynamic power management. *Computers, IEEE Transactions on*, 59(12):1625–1639, Dec 2010.

[12] C. Mastroianni, M. Meo, and G. Papuzzo. Probabilistic consolidation of virtual machines in self-organizing cloud data centers. *Cloud Computing, IEEE Transactions on*, 1(2):215–228, July 2013.

[13] K. Maurya and R. Sinha. Energy conscious dynamic provisioning of virtual machinesusing adaptive migration thresholds in cloud data center. *International Journal of Computer Science and Mobile Computing*, 2(3):74–82, March 2013.

[14] R. Rojas. *Neural Networks: A Systematic Introduction*. Springer-Verlag New York, Inc., New York, NY, USA, 1996.

[15] D. E. Rumelhart, J. L. McClelland, and C. PDP Research Group, editors. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1: Foundations*. MIT Press, Cambridge, MA, USA, 1986.

[16] G. Tong, H. Jin, X. Xie, W. Cao, and P. Yuan. Measuring and analyzing cpu overhead of virtualization system. In *Services Computing Conference (APSCC), 2011 IEEE Asia-Pacific*, pages 243–250, Dec 2011.

Multimedia Streaming over Cognitive Radios

Pranvera Kortoçi
Student number: 505709
pranvera.kortoci@aalto.fi

Abstract

Mobile phones, tablets, and laptops have become widely-used wireless communication devices that cover most of Internet-related activities. These devices are deployed in various areas and aspects of daily life. As a result, mobile traffic is facing a high demand from end-users. Video now represents more than 50% of all mobile traffic, and is expected to increase further in the near future. In order to cope with the demand for Radio Frequency (RF) spectrum, current approaches to access this resource call for conceptual modifications. This paper provides an overview of traditional techniques that provide access to the RF spectrum, and investigates the cognitive radio (CR) technology as a more efficient methodology of RF access. In addition, it details emerging models that enable sharing of the scarce RF resources, thereby increasing the availability of spectrum resources to the users. Finally, the paper presents the current multimedia protocols and ways in which multimedia streaming leverages the CR concept.

KEYWORDS: cognitive radio, spectrum sensing, adaptive video streaming, dynamic spectrum access

1 Introduction

Typical applications for smart phones, tablets, and other mobile devices range from multimedia messaging and video telephony/conference to video streaming [19]. As a consequence, network operators predict an increase of 1000 times in mobile data usage [8]. The RF spectrum is insufficient to satisfy the constantly increasing demand from the end-users to access it. Specifically, the RF spectrum suitable for mobile communications ranges from 30 to 3000 MHz. There are particular frequency bands within this spectrum that are underutilized either in terms of frequency, geo-location, or time.

The literature provides related case studies. For instance, Figure 1 shows the occupancy of the 2.3-2.4 GHz frequency band over several days in the city of Turku, Finland [13]. In Finland, this band is allocated for amateur services and wireless cameras. During a 24h period, the percentage utilization amounts to less than 30%. Scenarios like this demand for improved models that provide efficient utilization of resources.

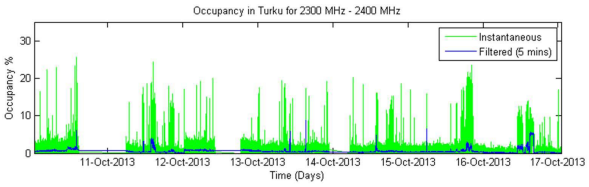

Figure 1: Percentage of bandwidth occupancy in Turku (10 - 17 Oct 2013) [13]

Both commercial and government clients share the RF spectrum. These clients access the spectrum either under a license or not. Clients that acquire a license to access a specific frequency band are provided with several rights, such as the availability of RF resources at any time and location, and guaranteed protection from interference. Other frequency bands require no license, hence allowing clients to access the unlicensed band in an opportunistic-basis. No protection rights are offered to these clients.

For instance, the Industrial, Scientific, and Medical (ISM) wireless frequency bands require no access license. The ISM bands comprise the 902-928 MHz, 2400-2483.5 MHz, and the 5725-5875 MHz band [21], which are internationally reserved for industrial, scientific, and medical applications. Devices operating in these bands must cause no harmful interference to licensees and are granted no interference-protection rights, thus no specific performance of such devices is guaranteed. In order for the ISM devices to avoid interference to licensees, the Federal Communications Commission (FCC) set rules and regulations in Part 15 of the Title 47 of the Code of Federal Regulations (47CFR) [10]. These rules consist of monitoring, and whereas needed, limiting transmission parameters such as maximum output power and spurious emission.

Several wireless communication technologies for wireless local and personal area networks operate in the ISM band. For instance, Bluetooth IEEE 802.15 and Wireless LAN IEEE 802.11 b/g operate in the 2.4 GHz ISM band, whereas Wireless LAN IEEE 802.11 a operates in the 5.8 GHz ISM band.

Static spectrum allocation policies assign frequency bands to licensees on a long-term basis covering a specific geo-location, leading to poor spectrum utilization [6]. The authors in [15] address the issue by introducing the Dynamic Shared Access (DSA) system, which allocates the RF spectrum under a licensed and unlicensed basis, as well as a combination of these. DSA suggests allocating part of the RF spectrum to secondary clients whenever a licensed frequency band is not employed. Users can opportunistically access these free portions of the spectrum, also known as *white spaces*. Moreover, mobile users equipment necessitate

to embed *cognitive-radio* capabilities in order to exploit dynamic access technologies. The term cognitive radio characterizes the ability of the CR user to explore and localize available spectrum portions, select the most appropriate ones, establish communication with other CR users to coordinate the access, and release the frequency portion if needed by a licensed user. More generally, CR defines "a radio that can change its transmitter parameters based on interaction with its environment" [11]. Figure 2 overviews the concept of a cognitive-radio system.

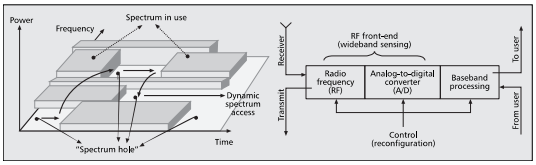

Figure 2: Overview of a *cognitive-radio* system [6]

The benefits of CRs consist of a higher utilization of the RF spectrum, increasing the responsiveness over time of current communication services, and promoting the interoperability of different communication systems in terms of transmission formats and operating frequency range. In detail, CRs incorporate several capabilities. Specifically, frequency adaptivity enables the CR user to dynamically select a more appropriate operating frequency in terms of performance. Moreover, modulation adaptivity modifies transmission characteristics such as the waveform in order to accommodate appropriate modulation. A CR can employ other capabilities such as transmission power adaptivity and geolocation of other CR users to increase the efficiency on accessing the RF spectrum and reduce interference to users. Figure 3 shows the CR network architecture.

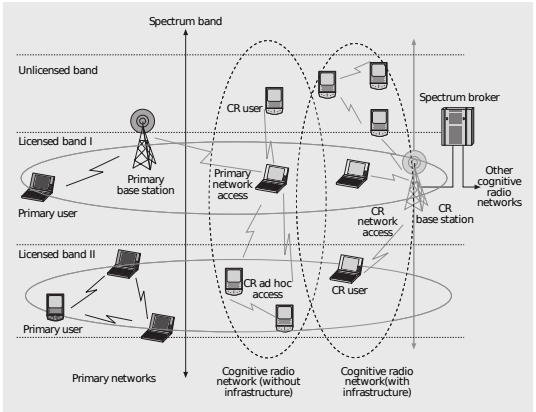

Figure 3: Cognitive and Primary network architecture [6]

The CR network and the active *primary network* cooperate. The primary network refers to the existing network infrastructure used by technologies such as 3G and 4G. CR users may access the licensed and unlicensed spectrum via their own base station, via an *ad hoc* CR access where CR users cooperate on both spectrum bands, and via the base station of the primary network. In the third, the CR user employs roaming functionalities within the Medium Access Control (MAC) protocol in order to profit from the primary network infrastructure (Figure 3). The CR network foresees the use of a spectrum broker, which performs scheduling among CR users that attempt to access the same available spectrum resources [14].

This paper gives a comprehensive understanding of spectrum sharing approaches and multimedia streaming protocols. To date, a crucial challenge of the multimedia delivery consists on addressing the intrinsic sensitive nature of the quality of service (QoS) to the available bandwidth. CR incorporates appealing capabilities to the media content delivery, thereby current studies in the field investigate the feasibility of embodying such capabilities into mobile devices. Achieving such interoperability would make the performance of the multimedia streaming reliable and efficient.

The rest of this paper is organized as follows. Section 2 introduces novel models of dynamic allocation that are supported by the existing infrastructure. Section 3 presents several multimedia streaming protocols, whereas Section 4 illustrates the multimedia streaming over cognitive radios. Finally, Section 5 concludes the paper.

2 Background

The FCC put considerable efforts to gaining a comprehensive understanding of the impact of efficiently using the RF spectrum. The FCC report [3] argues that "spectrum should be managed not by fragmenting it into ever more finely divided exclusive frequency assignments, but by specifying large frequency bands that can accommodate a wide variety of compatible uses and new technologies that are more efficient with larger blocks of spectrum". It recommends sharing the underutilized federal spectrum and implementing a new architecture that would enable the first shared-use of the spectrum. The major argumentation relates to the growth that the US economy would experience in terms of social benefits and the number of jobs created.

2.1 Spectrum Allocation and Sharing

Terms such as Authorized Shared Access (ASA), Licensed Shared Access (LSA), and Priority Access (PA) denote scenarios of sharing the RF spectrum among incumbent, secondary licensed, and priority-authorized users [15]. The licensed incumbent users are granted full interference protection rights and share the spectrum band with additional licensed users referred to as secondary in a mutually non-interfering policy. The priority-authorized user necessitates an authorization to access a specific frequency band over a specified time-interval.

The first attempt to pursue spectrum sharing comprises two tiers of users: an incumbent and a secondary licensed class of users. The model is called ASA/LSA and provides interference protection rights to both tiers. Although this model already provides sharing capabilities, it does not include the existence of an unlicensed user that would access the spectrum in an opportunistic manner, possibly increasing the spectrum access efficiency. Clearly, another class of

users with fewer rights could be included in a shared access model. This new class of unlicensed users is entitled to access the available RF spectrum at a certain time and/or geolocation subject to certain rules. These users are referred as General Authorized Access (GAA), and are not granted interference-protection rights.

The spectrum sharing aims at increasing the efficiency of accessing the RF spectrum. Thereby, the FCC recommendation report to the President of the US proposed to open the government spectrum band to commercial users by arguing for benefits it would have for economic development [3, 15]. FCC names the spectrum sharing among three different classes of users *Citizens Broadband Radio Service*, which requires a database that manages spectrum access. The database constantly updates the spectrum availability status, and the GAA users consult it prior to accessing the RF spectrum. The government spectrum is available partly to priority users under licensed form, and partly to general authorized users who might access the spectrum on an opportunistic basis. The federal and non-federal incumbent users would have interference protection rights to address the possible interference that Citizens Broadband Radio Service users might cause. Figure 4 shows rights and responsibilities of the three tiers as proposed by the FCC [1]:

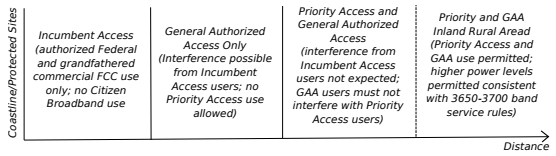

Figure 4: The three tiers rights and responsibilities [1]

Most of the current schemes for RF spectrum access are based on licenses. This is the case of network operators. They acquire a license which entitles them to operate exclusively at a certain frequency range. For instance, mobile operators guarantee to their clients the possibility of making a phone call at a certain time and location. Consequently, the novel models of accessing the RF spectrum must guarantee to these operators the allocation of such resources, in order to fulfill the clients demands at any time and circumstance. These users are classed as incumbent and are provided with interference protection rights. The operation within the shared access system of the PA users ensures interference protection rights from the GAA, and causing no harmful interference to the incumbent users. For example, mobile broadband operators could access the available spectrum under the PA model by requesting a priority access license for a specific frequency band. GAA users access the spectrum opportunistically whenever free spectrum is detected, and no interference to either incumbent or PA users is allowed. They have no interference protection rights from both, incumbent and PA users. WiFi hot spots service providers are GAA users with no interference protection rights [15]. Furthermore, the performance of the service is subject to the available spectrum band, the proximity of the incumbent and PA users, and the network conditions.

Unlicensed access to the TV band White Space (TVWS) is an example of the coexistence of incumbent and GAA users [9]. Television operators transmit over frequency ranges that are specified in their license. GAA users might access these bands under certain rules and limitations. Ultra Wide Band (UWB) falls under the GAA user class. This technique enables transmission over a wide range of frequencies at a low power. Users that implement this technique cause low interference, which is comparable to the noise floor of the licensee devices. This approach guarantees interference protection of the licensed devices.

In the following, we overview the ongoing work on implementing shared access protocols in Europe and in the US.

2.2 Standardization Efforts

The standardization of the shared access techniques is an ongoing process. Sharing access to the RF spectrum implies designing a network architecture that supports heterogeneous access mode, with appropriate tools to establish and cooperate such an access among different users. The Spectrum Access System (SAS) denotes such a network. The SAS is a reference point to clients attempting to access a frequency band. The FCC comprises all possible approaches to access the spectrum in the ASA, LSA, PA, and GAA models.

Nokia and Qualcomm were the first industry partners to propose the ASA concept. ASA aims at extending access to the International Mobile Telecommunications (IMT) band. IMT band converged national mobile networks bands into one, hence harmonizing the corresponding users to access a globally available frequency band. The user accesses the IMT band under a licensed form and is provided with an expected QoS.

The Radio Spectrum Policy Group (RSPG) further developed the ASA model in terms of accessible frequency band and allowed users, leading to the definition of an LSA model: “An individual licensed regime of a limited number of licensees in a frequency band, already allocated to one or more incumbent users, for which the additional users are allowed to use the spectrum (or part of the spectrum) in accordance with sharing rules included in the rights of use of spectrum granted to the licensees, thereby allowing all the licensees to provide a certain level of QoS” [16].

The major contribution of the LSA system is expanding the access beyond the IMT bands already foregone by the ASA model. Additional licensed users might access the band under regulatory policies that harmonize the coexistence of the initial and new users in terms of QoS and interference management. The LSA provides new users with QoS rights, hence it guarantees a predictable QoS [16]. These users are subject to the regulations of national authorities, therefore they might be granted or not same rights as initial users.

The PA model identifies an access scheme for specific users located in determined geographic areas with distinct requirements in terms of QoS to perform their activities. The PA class includes users such as hospitals, utilities, and government institutions [1].

An overview on the ongoing standardization process in EU and USA follows.

2.2.1 EU efforts toward ASA/LSA in the 2.3 GHz band

In 2013, the RSPG revised the definition of LSA as follows: "A regulatory approach aiming to facilitate the introduction of radiocommunication systems operated by a limited number of licensees under an individual licensing regime in a frequency band already assigned or expected to be assigned to one or more incumbent users. Under the Licensed Shared Access (LSA) approach, the additional users are authorised to use the spectrum (or part of the spectrum) in accordance with sharing rules included in their rights of use of spectrum, thereby allowing all the authorized users, including incumbents, to provide a certain Quality of Service (QoS)" [17].

The 2.3-2.4 GHz band in Europe is a good example of a poorly utilized band. The World Radio Conference assigned this band to the IMT in 1997. However, Europe poorly utilizes this band due to the high activity of incumbent users. Consequently, Europe mostly focused on defining the ASA/LSA model in the 2.3-2.4 GHz RF band. In [2], the European Telecommunication Standard Institute (ETSI), after starting a standardization process for ASA/LSA, defined in its report the technical and operational conditions to implement such model in this band. The ASA/LSA model that RSPP and ETSI defined for the 2.3-2.4 GHz band could extend to the 1.7 GHz and the 3.5 GHz bands.

In Europe, Finland was the first country to realize spectrum sharing under the ASA/LSA model in an LTE network operating in the 2.3 GHz band in April 2013. The United Kingdom, as the rest of Europe, defined a shared access method within underutilized bands. In November 2010, UK published the *Implementing Geo-Location* consultation. Ofcom became part of the statement regarding the White Space Devices (WSD), which was published in September 2011. WSD devices are allowed to access the TV White Space as long as no harmful interference is caused to existing operators in the TVWS band. Ofcom explicitly contributed in defining a geo-location database which would monitor and control the emission levels of the WSD devices within the TVWS band, by limiting or ceasing transmissions if required [16]. Trials were performed in Cambridge, aiming at supplementing services such as wireless broadband and attempt to fulfill the high demand for mobile data. The tested WSD devices support Wi-Fi, rural broadband, and machine-to-machine communications (M2M) services. Moreover, UK explicitly proposed merging the UK and the European shared access methodology into one, arguing the benefits to both citizens and consumers.

2.2.2 US 3.5 GHz shared band under ASA/LSA

The US efforts concentrate on enabling commercial users to access the government bands, such as the federal radar system band. This is the case of the 3550-3650 MHz, which is known as the 3.5 GHz Small Cells band. The FCC proposed the 3.5 GHz band as suitable to be exploited by *small cells*, also called femto cells. The base stations of these cells operate at low power, and are typically employed to extend wireless coverage to small indoor or outdoor areas.
The FCC report [1] overviews the 3.5 GHz band characteristics. In the US this band is allocated to the Radiolocation Service (RLS) and to the Aeronautical Radionavigation Service (ARNS) for federal use on a primary basis. The advantage of deploying small cells consists of reducing the risk of interference, which derives from a smaller coverage area of small cells with regard to macrocells. Consequently, the frequency reuse would increase along with the network capacity. Moreover, small cells respond greatly to the spectrum sharing system on a geographical basis. In fact, they operate at low power and limited signal propagation, thus enabling spectrum sharing with predicted QoS.

The small cell concept fits well within a 3-Tiers architecture accessing the 3.5 GHz federal band (Figure 4). The FCC framework encloses directives that would enable the implementation of a SAS scheme for the three Tiers. Moreover, it suggests investigating and implementing mitigation techniques in order to reduce the interference risk factor.

3 Multimedia Streaming Protocols

Multimedia refers to a variety of content such as audio, video, animation, and text. Multimedia streaming is regulated by protocols that define entities of such a system, and ways of exchanging information among them in order to provide multimedia content to the users. These protocols experience continuous conceptual variations. Current protocols aim at harmonizing among a wide range of device capabilities and multimedia streaming features. Currently, end-users demand higher quality of experience (QoE) over a wide range of transmission systems, such as unicast (Video on Demand delivery) and multicast (mobile TV delivery), with QoE being a performance indicator of the service experienced by the user.

Multimedia streaming is a power-hungry application that causes short battery-life on mobile devices. The work in [12] investigates streaming techniques from a mobile device perspective. It explores the power-saving mechanisms that wireless network interfaces such as Wi-Fi, WCDMA/HSPA, and LTE employ. These technologies adopt various streaming techniques such as: encoding rate streaming where the streaming rate equals the encoding one; rate throttling where the streaming rate is higher than the encoding rate; buffer-adaptive streaming which delivers multimedia content solely when the playback buffer drains to a certain threshold; fast caching in which the user downloads the entire video content when it first connects to the server; and rate adaptive streaming over HTTP where the streaming rate adapts to the available channel bandwidth. The findings in [12] claim that servers mainly adopt fast caching and throttling, while video players favor encoding rate and buffer adaptive mechanisms; the network operator has no role in choosing the streaming technique because the wireless interface plays no role on the decision; low quality videos provide a shorter joining time (initial playback delay), thus Wi-Fi offers a shorter delay than HSPA or LTE; the video quality does not significantly affect the energy consumption; and fast-caching and throttling reduce the energy consumption in an uninterrupted streaming scenario.

In the following, we detail the major multimedia streaming protocols and methodologies to cope with the bandwidth limitations.

3.1 Dynamic Adaptive Streaming over HTTP

The Third Generation Partnership Project (3GPP) specifies in its draft version the Progressive Download over HTTP and the Dynamic Adaptive Download over HTTP (DASH) protocols [4].

Traditional streaming techniques employ protocols that keep track of the state of the client. For example, the client and the streaming server are connected while the streaming server delivers the multimedia content to the client as a continuous stream of packets [20]. Alternatively, clients may download the media content via the HTTP protocol, and play it from a local storage. Thus, it is possible to rely on the existing HTTP servers and caches to store the media content. Progressive Download over HTTP enables the user to sequentially download the media content as a media file that consists of continuous media, therefore the user can start playing the content prior to the full download [5]. The client uses the HTTP GET or partial GET commands to request a certain media file and obtain it. Progressive download over HTTP treats the media file as a continuous stream, thus leading to poor performance. Moreover, the interruption of the streaming of the media content would result in resource wastage because the progressive download has already started and the media content is stored in the HTTP servers/caches. Furthermore, progressive download over HTTP does not support bit rate adaptivity and live media services [20].

The DASH protocol overcomes the resource wastage and the poor performance in terms of adaptivity of the Progressive Download over HTTP protocol. Unlike its predecessors such as Real-Time Transport Protocol (UDP/RTP), DASH stands for a stateless server architecture, moving the control logic from the server to the client. Figure 5 shows the 3GPP-DASH architecture specified in [5]. The 3GPP-DASH client establishes a HTTP-URL connection with the server in order to provide streaming data to the user. The Media Presentation Description (MPD) provides the client with metadata required to enable such a connection. The MPD is a XML document containing information regarding the multimedia content on the server and the corresponding HTTP-URLs.

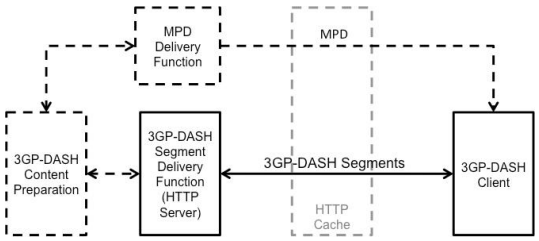

Figure 5: DASH Protocol Architecture [5]

The architecture in Figure 5 also shows a 3GPP-DASH client accessing the MPD. The 3GPP-DASH client requests segments of multimedia data from a HTTP server by exploiting the MPD metadata. The delivery of the multimedia content employs HTTP/1.1 delivery protocol.

The Progressive Download and the DASH protocols support QoE reports. These reports are optional and may be triggered via the MPD. The quality metrics supported by both Progressive Download and DASH [5] are the list of HTTP request-response transactions, average throughput, initial playout delay, and buffer level, while DASH supports two more metrics, such as the list of representation switch events and the MPD information. The HTTP request-response transaction metric evaluates whether the server, based on the channel conditions, fulfills the client request to download multimedia content at a certain bit rate and quality. The average throughput estimates the channel capabilities to cope with the clients demands, whereas the initial playout delay is the time interval between the request and the playout of the media content. DASH retrieves such metrics and maps them into a *QoE map*. 3GPP-DASH also offers on demand, live, and time-shift services [5, 20].

3.2 Scalable Video Coding

Scalable Video Coding (SVC) is an extension of the H.264/AVC video coding standard. The H.264/AVC conceptual blocks comprise the Video Coding Layer (VCL) and the Network Abstraction Layer (NAL), where the first provides an encoded version of the source content and the second takes over the encoded version of the source. The NAL formats and provides header information to the VCL data, hence enabling it to be employed in various channel conditions scenarios.

H.264/AVC organizes a picture into coding units, *macroblocks* and *slices*. Each picture consists of small rectangular pictures called macroblocks, and each macroblock consists of slices. These slices host either spatial or temporal prediction information to be employed by the macroblocks. H.264/AVC supports three slice coding types: I-slice/Intra-picture with spatial prediction from adjacent slices; P-slice with one-directional prediction from Intra and Inter-pictures, where Inter-pictures refer to adjacent pictures for encoding; and B-slice with bidirectional prediction. The H.264/AVC capabilities trade-off for high decoder complexity [19]. Therefore, SVC addresses the decoding complexity of such protocol, and increases the scalability degree of the H.264/AVC.

The key feature of the SVC consists on dividing the source bit stream into portions. Instead of encoding and transmitting the entire bit stream at a specific rate, these portions are encoded at several temporal or spatial resolutions, hence enabling users to request for the more appropriate encoded version [19]. The rate of the singular partitions is lower than that of the source bit stream. However, their reconstruction leads to a higher rate than the original bit stream exhibits.

SVC encodes the video stream into Base Layers (BL) and Enhancement Layers (EL). These layers can be stored independently and the encoded versions of the EL are made available for download at a latter moment, whereas the channel throughput allows. By dividing the original bit stream into portions, it increases the degree of intervention and adaptation that can be applied by leveraging temporal, spatial, and quality scalability.

- Spatial scalability. The picture size of the substreams varies. Each substream/layer supports a spatial reso-

lution. Each spatial layer employs intra and motion-compensated prediction.

- Temporal scalability. The frame rate of the substreams varies. The substream consists of one temporal base layer, and one or more ELs. The EL is obtained by encoding B-slices. Figure 6 shows ways of enabling temporal scalability. In Figure 6 (a), there are two reference pictures denoted as 0 and 1, corresponding to temporally preceding and succeeding pictures. T_k represents a set of temporal layers. The temporal layer of the two reference pictures is T_0, which is lower than the temporal layer identifier of the predicted pictures. Therefore, a temporal layer can be encoded independently of temporal layers whose layer identifier $T > k$. The EL pictures are retrieved by encoding B-pictures in a hierarchical prediction structure.

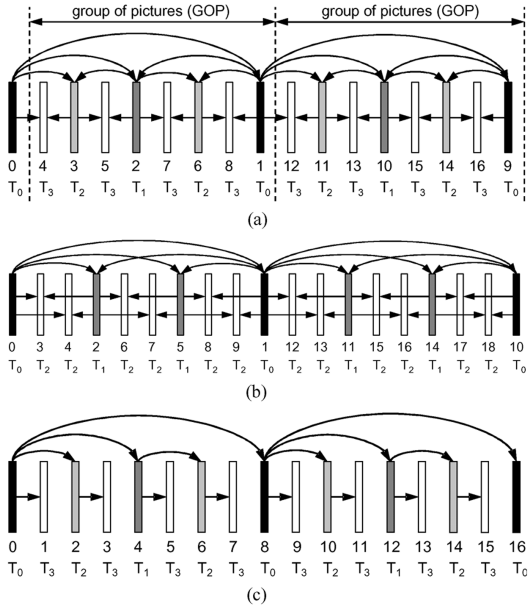

Figure 6: Temporal scalability: coding with hierarchical B-pictures [19]

- Quality scalability. The substreams provide the same spatio-temporal scalability as the original one at lower fidelity, intended as the Signal to Noise Ratio (SNR) of the bit streams.

The term *scalability* represents the capability of partitioning a bit stream and performing several operations on these partitions. The data contained in each of the partitions is lower than the theoretical proportion, however each partition data inherits crucial information in order to ensure that the partitions could merge back together and form a new stream. This stream has lower quality than the original one. The quality of the substreams though, is *relatively* higher than the original one. As the data contained in the substreams is lower than the original one, a relative comparison implies that the quality of the substreams is higher.

SVC encodes only one bit stream at the highest picture size and bit rate, and then it applies the scalability in order to obtain numerous versions of bit streams with different resolution or bit rate, hence enabling adaptation with regard to the end-user device and channel capabilities.

3.3 Improved Dynamic Adaptive Streaming over HTTP (iDASH)

DASH provides clients with the most appropriate video chunk for a specific end-user device capability and network conditions, whereas SVC encodes the multimedia content into BL and ELs, thus enabling the users to improve the QoE by downloading ELs whenever the channel conditions allow it. The authors in [18] merged the concept of the SVC into the DASH technique. SVC adds one degree of adaptation to the DASH technique structuring the media content into two layers, one corresponding to the most appropriate chunk in terms of bit rate at a specific time, and the other corresponding to the most appropriate representation of such a chunk in terms of quality for a specific channel condition. Therefore, the client prioritizes between multiple representations of the same data chunk. The iDASH client sends several HTTP GET requests within a time interval that correspond to the base and enhancement layers. However, if the channel conditions allow no ELs to be downloaded, the corresponding request is omitted. Therefore, iDASH achieves higher responsiveness and quality adaptation in scenarios of fast variations of the channel conditions over time.

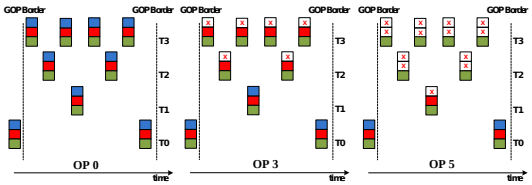

Figure 7: iDASH: SVC encoding of temporal and quality layers [18]

iDASH offers multiple Operation Points (OP) within a bit stream. An OP defines an SVC sub-stream at a specific bit rate and quality level. SVC allows for on-the-fly adaptation, adapting the number of layers of the sub-stream to the channel conditions.

Figure 7 shows the concept of the OPs. A sub-stream consists of a base layer (bottom), and two enhancement layers Q1 (center) and Q2 (top). In order to obtain several OPs, starting from the higher temporal levels (T_n), the higher quality layers are dropped in turns of Q2 and then Q1. Consequently, the bit rate decreases allowing the user to smoothly adapt to the channel conditions.

The iDASH technique outperforms plain DASH in terms of responsiveness over time to the bit rate and congestion control. iDASH achieves the latter by gradually downgrading the request for higher ELs to adapt the bit rate of the sub-stream to the ongoing congestion rate.

The authors in [7] investigate ways of selecting a proper quality of multimedia content for DASH with SVC scenario

under both, constant and variable download rate. DASH splits a media content into several quality profiles, whereas SVC splits the raw media content into subsets of bit-stream. These bit-streams can further be split into segments. Merging DASH with SVC allows the user to download segments before the playback deadline and increase the quality.

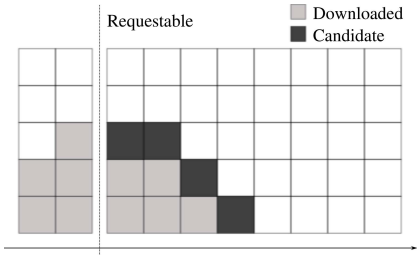

Figure 8: Decision policy for a client in a DASH with SVC scenario [7]

Figure 8 shows how the user chooses to either increase the quality of the current segment by downloading in vertical, or download in horizontal for future segments. The first decision denotes a backfilling approach, whereas the second denotes prefetching. The user adapts its downloading policy to a diagonal one, which is influenced by the variations of the download rate. [7] argues that the slope becomes flatter under low, highly variable, and highly persistence download rate. Predicting the available bandwidth would allow users to accordingly modify the slope of the decision policy, thus experiencing a flatter and higher QoE.

4 Multimedia Streaming over CRs

So far we have introduced cognitive radios and multimedia streaming independently. We now present multimedia streaming specifically targeted to cognitive radios.

Multimedia streaming suffers from the scarce availability of bandwidth, thereby it exposes the user to a highly sensitive QoE. The multimedia delivery over CR networks consists of sensing a spectrum band and opportunistically accessing the detected free spectrum, and thereby increasing the available resources. Consequently, the users experience a more stable QoE over time.

4.1 SVC for 3-Tiered spectrum sharing

The multimedia streaming QoE is subject to the available bandwidth allocated to deliver such service. As mentioned in Section 1, the RF spectrum is a common source for several entities with different spectrum rights. Figure 9 shows a three tiers model sharing the white space to deliver adaptive video streaming [22].
The three tiers are provided with different rights. Tier-1 and Tier-2 are licensed users of the spectrum band, and thereby guaranteed interference protection rights and available bandwidth for access. The unlicensed Tier-3 opportunistically accesses the white space on an access-by-rules approach. T3 users are lower in priority than T1 and T2, hence they

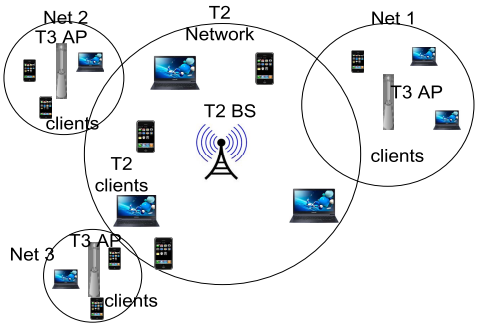

Figure 9: Spectrum sharing among three tiers of users [22]

provide adaptation in terms of transmission power to avoid harmful interference to the licensees. Such adaptation shapes the coverage area of the T3 network. The closer the licensed users, the lower the power transmitted by the T3 network. Embedding the T2 mobility pattern information into the T3 adaptation mechanism leads T3 users to respond accordingly by decreasing the transmission power, thus reducing the overlap area of the T2 and T3 Access Points (APs) as shown by Figure 9.

As mentioned in Section 3.1, DASH consists of progressively requesting portions of the video stream (segments) of various resolutions, thus it adapts to the channel conditions by following its average throughput, which is estimated based on previous observations of the latter. DASH adaptation performs satisfactorily when channel conditions are stable over time; however, it performs poorly under fast-variable channel throughput. Partial information about the mobility of T2 users would enable the estimation of the channel throughput, though DASH lacks of mechanisms to exploit this information and perform such prior evaluation.

In contrast, SVC already provides on-the-fly adaptation of the bit rate and QoE. Embedding information about mobility pattern into SVC would increase the adaptation degree, allowing the user to smoothly downgrade the bit rate and reduce QoE fluctuations. The SVC client chooses the more appropriate encoded version of the bit stream for download. The ELs download adapts to the channel throughput estimation. Furthermore, T3 users manage to adapt the transmission power accordingly.

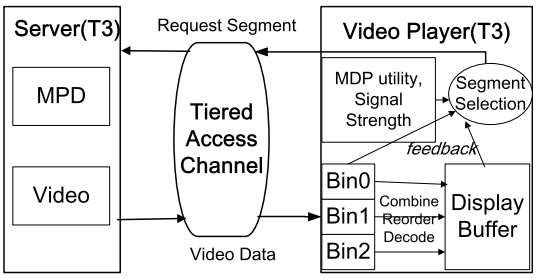

Figure 10: SVC within a whitespace scenario [22]

The system shown in Figure 10 merges the concept of the T2 mobility prediction in a white space environment into the

SVC protocol [22]. It consists of three basic components: the video server, the channel access, and the video player. The video server stores SVC video formats along with the corresponding MPD files. On the client side, the video player implements and runs an algorithm in order to request the appropriate video segment in terms of its resolution. The algorithm evaluates the T3 user throughput based on the SNR measured signal of the T2. Consequently, the video player sends a request to the video server to download a specific version of a segment via the tiered access channel. The decision is based on the T2 SNR, and the prior and predicted white space channel throughput. The white space channel history updates each time a segment request is sent, and the collected information is embedded back into the algorithm in a loop basis.

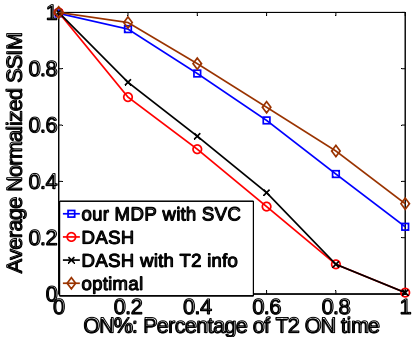

Figure 11: Video quality as a function of the time the T2 user is ON [22]

Figure 11 shows the video quality as a function of the time the T2 user is ON. The video quality is expressed by the Structural SIMilarity (SSIM) index which measures the similarity between two images, where one of the images provides the perfect quality. The SSIM takes decimal values in the -1 to +1 range, where +1 represents the case of two identical images. SSIM is then normalized to the average SSIM of the experiment video data. Authors in [22] compare the performance of four algorithms: the optimal algorithm assumes that the throughput of the white space channel is known throughout the entire video duration; DASH algorithm adapts to the current throughput variations, although it has no means of exploiting predicted throughput variations; DASH with T2 info constitutes a modified version of conventional DASH embedding information regarding the ON/OFF state of T2 over the next τ seconds of the video data; and MDP with SVC algorithm employs a two-state Markov chain to represent the presence of the T2 user, either ON or OFF. Furthermore, the duration of the T2 user in either states is modeled as a geometric distribution. The algorithm then, based on a Markov Decision Process (MDP), selects the more appropriate video segment and quality level to request for download.

Figure 11 shows that the video quality decreases considerably as the time period of T2 ON increases. The performance of DASH and DASH with T2 info algorithms degrades faster than the MDP with SVC, which closely follows the optimal one. In the worst case of the T2 user being in the state ON

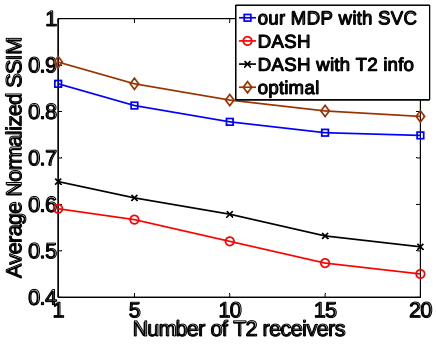

Figure 12: Video quality as a function of the density of T2 users in ON state [22]

for the entire duration of the video data, MDP with SVC achieves higher SSIM by downloading enhancement layers whenever possible. Figure12 however, shows that the video quality is less reactive to the variation of the number of T2 users in the ON state. The gap between the optimal and the MDP with SVC performance keeps constant and is relatively small compared to the gap between the optimal and DASH algorithm.

5 Conclusion

Multimedia streaming over CRs consists of sensing and accessing free portions of the RF spectrum to increase the available bandwidth for multimedia delivery. Implementing CRs capabilities into mobile devices would enable current streaming techniques to leverage features such as the RF utilization prediction and feedback data regarding the history of the channel condition. Merging them together would lead to higher responsiveness and adaptation of the multimedia streaming quality over time and RF utilization.

In order to exploit the white space, end-user devices must provide real-time signal processing capabilities that would expose these devices to high energy consumption, thus limiting their battery-life. Current research in the field investigates the feasibility of transparently embedding CR capabilities into mobile devices. Research aims at enabling such capabilities with no significant impact on the battery-life of mobiles devices, while increasing the QoE of multimedia streaming.

References

[1] Amendment of the commissions rules with regard to commercial operations in the 3550-3650 MHz Band. "`http://apps.fcc.gov/ecfs/document/view?id=7022080889`", December 2012.

[2] ETSI reconfigurable radio systems workshop. http://www.etsi.org/news-events/news/410-reconfigurable-radio-systems-workshop-12-december-2012-sophia-antipolis-france, December 2012.

[3] Report to the president realizing the full potential of government-held spectrum to spur economic growth. "http://www.whitehouse.gov/sites/default/files/microsites/ostp/pcast_spectrum_report_final_july_20_2012.pdf", July 2012.

[4] 3GPP. Transparent end-to-end Packet-switched Streaming Service (PSS); Progressive Download and Dynamic Adaptive Streaming over HTTP (3GP-DASH). "http://www.3gpp.org/DynaReport/26247.htm".

[5] 3GPP. 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Transparent end-to-end Packet-switched Streaming Service (PSS); Progressive Download and Dynamic Adaptive Streaming over HTTP (3GP-DASH) (Release 10). http://www.3gpp.org/DynaReport/26247.htm, December 2014.

[6] I. F. Akyildiz, W.-Y. Lee, M. C. Vuran, and S. Mohanty. A survey on spectrum management in cognitive radio networks. *Communications Magazine, IEEE*, 46(4):40–48, 2008.

[7] T. Andelin, V. Chetty, D. Harbaugh, S. Warnick, and D. Zappala. Quality selection for dynamic adaptive streaming over http with scalable video coding. In *Proceedings of the 3rd Multimedia Systems Conference*, pages 149–154. ACM, 2012.

[8] T. Cisco. Cisco visual networking index: global mobile data traffic forecast update, 2012–2017. *Cisco Public Information*, 2013.

[9] Coops and Martin. First report and order, in the matter of revision of part 15 of the commission's rules regarding ultra-wideband transmission systems. Technical report, Federal Communications Commission, April 2002. https://apps.fcc.gov/edocs_public/attachmatch/FCC-02-48A1.pdf.

[10] FCC. Rules and Regulations for Title 47. "http://www.fcc.gov/encyclopedia/rules-regulations-title-47".

[11] FCC. Notice of proposed rule making and order. "https://apps.fcc.gov/edocs_public/attachmatch/FCC-03-322A1.pdf", December 2003.

[12] M. A. Hoque, M. Siekkinen, J. K. Nurminen, M. Aalto, and S. Tarkoma. Mobile multimedia streaming techniques: QoE and energy saving perspective. *Pervasive and Mobile Computing*, 16:96–114, 2015.

[13] M. Hoyhtya, M. Matinmikko, X. Chen, J. Hallio, J. Auranen, R. Ekman, J. Roning, J. Engelberg, J. Kalliovaaras, T. Taher, et al. Measurements and analysis of spectrum occupancy in the 2.3–2.4 GHz band in Finland and Chicago. In *Cognitive Radio Oriented Wireless Networks and Communications (CROWNCOM), 2014 9th International Conference on*, pages 95–101. IEEE, 2014.

[14] O. Ileri, D. Samardzija, and N. B. Mandayam. Demand responsive pricing and competitive spectrum allocation via a spectrum server. In *New Frontiers in Dynamic Spectrum Access Networks, 2005. DySPAN 2005. 2005 First IEEE International Symposium on*, pages 194–202. IEEE, 2005.

[15] W. Lehr. Toward more efficient spectrum management, new models for protected shared access. Technical report, MIT Communications Future Program. http://cfp.mit.edu/publications/CFP_Papers/CFP%20Spectrum%20Sharing%20Paper%202014.pdf.

[16] Radio Spectrum Policy Group. Report on Collective Use of Spectrum (CUS) and other spectrum sharing approaches. "http://rspg-spectrum.eu/_documents/documents/meeting/rspg26/rspg11_392_report_CUS_other_approaches_final.pdf", November 2011.

[17] Radio Spectrum Policy Group. RSPG Opinion on Licensed Shared Access. "https://circabc.europa.eu", November 2013.

[18] Y. Sánchez de la Fuente, T. Schierl, C. Hellge, T. Wiegand, D. Hong, D. De Vleeschauwer, W. Van Leekwijck, and Y. Le Louédec. iDASH: improved dynamic adaptive streaming over HTTP using scalable video coding. In *Proceedings of the second annual ACM conference on Multimedia systems*, pages 257–264. ACM, 2011.

[19] H. Schwarz, D. Marpe, and T. Wiegand. Overview of the scalable video coding extension of the H. 264/AVC standard. *Circuits and Systems for Video Technology, IEEE Transactions on*, 17(9):1103–1120, 2007.

[20] T. Stockhammer. Dynamic adaptive streaming over HTTP–: standards and design principles. In *Proceedings of the second annual ACM conference on Multimedia systems*, pages 133–144. ACM, 2011.

[21] U.S. Government Publishing Office. Operation within the bands 902-928 MHz, 2400-2483.5 MHz, and 5725-5850 MHz. "http://www.ecfr.gov/cgi-bin/text-idx?SID=dde1f198f15744f94c43c74c703b773f&node=se47.1.15_1247&rgn=div8".

[22] X. Wang, J. Chen, A. Dutta, and M. Chiang. Adaptive Video Streaming over Whitespace: SVC for 3-Tiered Spectrum Sharing. In *IEEE INFOCOM*, 2015.

IPv6 over networks of resource-constrained nodes

Lauri Luotola
Student number: 84703B
lauri.luotola@aalto.fi

Abstract

This paper describes the key technologies and challenges involved in transmitting IPv6 packets over Low-Power Wireless Personal Area Networks (LoWPANs). LOWPANs are in use especially in sensor networks that are being used for applications such as environmental monitoring and home automation. Bringing IPv6 to these networks is essential as the amount of devices can be expected to grow exponentially. Various implementions of LOWPANs exist, of which ZigBee, NFC and Bluetooth Low Energy are examined and compared. Majority of these technologies are still being actively developed and not all of them are direct competitors, but instead have their own use cases.

KEYWORDS: IoT, 6lo, IPv6, 6LoWPAN, BLE, ZigBee, NFC, IEEE 802.15.4

1 Introduction

The Internet of Things (IoT) is about extending the scope of the Internet to cover physical objects. Wireless nodes are expected to outnumber conventional computer nodes that we think of as the Internet of today. These wireless devices, often referred to as sensors, are typically strictly resource constrained and must operate unattended for long periods of time, hence it is essential for them to transmit their data to the Internet efficiently, with low power consumption. IoT can be applied to numerous fields and offer wide-ranging possibilities; applications have been deployed in medicine, agriculture, environmental matters, military, toys and many others. Typical applications include devices such as health monitors, environmental sensing and proximity sensors. [10] These devices may be deployed in the magnitude of thousands, thus cost savings quickly add up and they are aimed to be produced at low cost.

The IP (Internet Protocol) architecture was arguably not designed for resource-constrained devices, having its origins in the 1970s for connecting general purpose computers using wired networking technologies such as Ethernet. IPv4, a version of the Internet Protocol, has been widely and very successfully deployed on hundreds of millions of hosts and routers in private and public networks alike. Being initially designed in 1982, this growth rate is remarkable. The Internet Engineering Task Force (IETF), responsible for the standarding efforts of the IP protocol suite, identified the need for a new revision to address the problems brought by the huge success of the protocol. This led to the specification of IPv6 in 1998. [10] IPv6 is an evolution of IPv4 with no change in the fundamental and architectural principles of the IP protocol suite. The reason IPv4 is still prevalent is mainly the cost and complexity of migration, which has made the adoption rate rather slow.

A Low-power Wireless Personal Area Network (LoWPAN) is a simple low-cost communication network consisting of devices conforming to the IEEE 802.15.4 standard. They are typically used to connect resource-constrained devices such as wireless sensors. Characteristics of LoWPANs include small packet size, low bandwidth, low-cost and low reliability of the connected devices. The devices often tend to sleep for longer periods of time in order to save energy and the connection may be unreliable. [5, 10]

There is a working group for IETF that aims to adapt IPv6 for IoT devices, namely 6LoWPAN, which also refers to the effort itself [6]. Prior to 6LoWPAN, many vendors embraced proprietary protocols because the IP architecture was thought to be too resource-intensive to be operated on such devices [10]. A typical smart object has very limited memory, not enough to operate the existing implementations of the IP protocol family; for this reason a number of non-IP stacks were developed [10]. However, the field has since matured and 6LoWPAN brings significant improvements over older protocols, introducing an adaptation layer that enables efficient IPv6 communication over various links, such as IEEE 802.15.4 or Bluetooth Low Energy. [5, 7] This paper discusses these different applications and their differences.

2 The advantages of IPv6

IPv6 aims to overcome the limitations of IPv4 and eventally replace it, thus enabling the Internet to scale further. One of the most important reasons for this revision is to overcome the problem of exhausting address space. IPv6 expands the IP address space from 32 to 128 bits - and as a result of this, increases the required maximum transmission unit (MTU) from 576 to 1280 bits compared to IPv4. Even though the benefits of IPv6 are clear, the majority of Internet traffic is still handled by IPv4.

Considering the potential for exponential growth in the number of Internet of Things applications, and connected devices in general, IPv6 is an ideal protocol due to the large address space it provides. In addition, IPv6 provides tools for stateless address autoconfiguration; the limited processing power and large number of devices in a LoWPAN network make automatic network configuration and statelessness highly desirable.

2.1 IPv6 over Low-power Wireless Personal Area Networks

Mapping the IPv6 network to the IEEE 802.15.4 network presents several design challenges - IEEE 802.15.4 devices are characterized by short range, low bit rate, low power, and low cost. IP protocols generally assume that the link is always on and constantly listening for packets. However, this behaviour of idle listening is not suitable in context of the low-power requirement for IoT devices. Also, many of the devices employing IEEE 802.15.4 radios will have limited computational power, memory, and energy availability. IEEE 802.15.4 networks are usually ad-hoc networks since their location is usually not predetermined. The connection is usually unreliable, especially when compared to wired links such as Ethernet. IEEE 802.15.4 has a maximum data rate of 250,000 bits/s and a maximum output power of 1 mW. The devices have a nominal range on the order of a few tens of meters [10].

Devices within LoWPANs are expected to be deployed in large numbers and to have limited capabilities. The large number of deployed devices in the network requires scalable technologies, raising the need for a large address space, which however is well met by IPv6. Due to the devices often being used as sensors, their location may be hard to reach. This poses the need for automatic configuration and management, thus the implementations of LoWPAN should preferably work with minimal configuration and be able to automatically adjust to network problems. IPv6 provides tools for stateless address autoconfiguration, which is particularly suitable for sensor network applications and nodes with limited processing power.

The common underlying goal is to also reduce processing requirements, bandwidth and power consumption, therefore packet overhead should be minimized. Given the limited packet size, headers for IPv6 and layers above must be compressed whenever possible.

Another challenge IPv6 faces is the required MTU of 1280 bits — low-power radio links typically do not support such a large payload; IEEE 802.15.4 frame only supports a payload of 127 bits. To overcome this, there has to be an adaptation layer to allow the transmission of IPv6 datagrams. The MTU size of IEEE 802.15.4 is purposely small to cope with limited buffering capabilities and to limit the packet error rate since the bit error rate can be relatively high. Fragmentation of IEEE 802.15.4 frames may be required at the 6LoWPAN adaptation layer when the IPv6 payload exceeds the MTU. [10]

2.2 Network topologies

IPv6 over LoWPANs has been implemented using various technologies that utilize different topologies for node-to-node communication. The IEEE 802.15.4 standard [6] defines two types of devices: full function devices (FDD) and reduced function devices (RFD). FDDs can act as network coordinators and function in any topology, whereas RFDs are limited to star topology, communicate only with the network coordanators and can be applied to resource-constrained devices. In most cases, the RFDs are battery powered; therefore, in order to prolong the lifetime of the network, it is essential to minimize the power consumption in the communication between nodes without compromising network connectivity. This has implications on the routing protocols because the shortest path may not always be the most energy efficient.

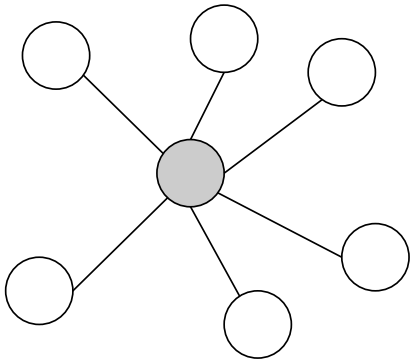

Figure 1: Star topology is the only possibility if the sensors never have the radio on for listening. Nodes only communicate with a central router.

In star topology, every node is connected to a central router, forming a graph with the topology of a star, as can be seen from Figure 1. Star topologies include provisioning only a subset of devices with packet forwarding functionality. If these devices use various kinds of network interfaces, the goal is to seamlessly integrate the networks built over those different technologies. Star networks are the only types of networks possible if the devices never have the radio on to listen for transmissions from neighbors. [10] Star topology is simple and useful, but constrains the range of the netowrk to that of the physical transmission range of the transceivers.

To allow for the network range to extend beyond this, individual nodes must be able to receive transmissions from each other. Figure 2 is an example of mesh network topology. A mesh network is a network topology in which each node relays data for the network. All nodes cooperate in the distribution of data in the network. Mesh topologies imply multi-hop routing to a desired destination. Intermediate devices act as packet forwarders at the link layer, similarly to routers at the network layer. Mesh topology also provides added reliability as it can construct redundant paths through the network - if a node goes down, it can reroute the network traffic through other nodes. To be able to form mesh networks, the radio transceivers of the individual nodes have to be turned on periodically to listen for neighbours' communication. [10]

A tree topology connects multiple star networks to other star networks. If the connection between each of the star topology networks fails, those networks would be unable to communicate with each other. However, computers on the

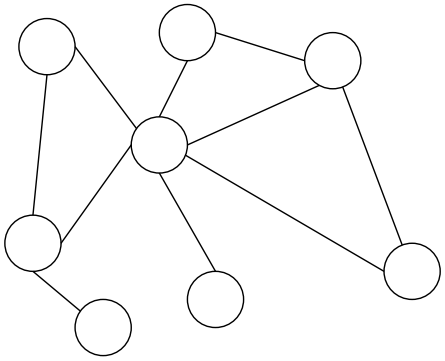

Figure 2: In a mesh topology, all the nodes can talk with each other, allowing the network to dynamically extend and increase redundancy.

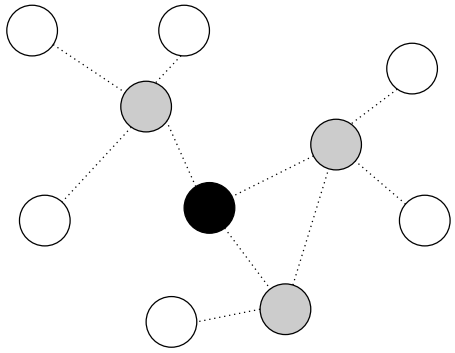

Figure 3: An example of a ZigBee network. The black node is the coordinator, gray ones are routers and the white ones are end devices.

same star topology would still be able to communicate with each other.

3 Implementations of low-power networking

In the late 1990s, there was a strong movement toward defining a new network architecture, designed to provide a standards-based protocol for interoperability of sensor networks, which was named ZigBee [10]. The design focused on control applications such as home automation over a low-power wireless communication medium. ZigBee initially defined its own networking stack that was incompatible with existing network standards such as IP, however later it moved towards adopting IP as its communication mechanism. The current version of ZigBee is a networking layer built on top of the IEEE 802.15.4 standard, which defines the physical and MAC layers. The ZigBee standard has existed since 2004, and it was implemented with different profiles for each operating environment. [1]

That saved manufacturers from having to implement a broad set of capabilities, but also means there are separate networks for different uses of ZigBee in many cases. ZigBee identifies three node types: end-device, router and coordinator. A coordinator is an FDD that manages the whole network, routers have routing capabilities and the end-devices act as peers that only transmit data. ZigBee supports star, tree and mesh topologies. A benefit of ZigBee is that its nodes can remain in sleep mode most of the time, thus extending battery life.[9, 1] Figure 3 shows an example of a ZigBee network.

Z-Wave is an alliance that developed its own patented low-power RF technology for home automation. The technology is not IP-based and thus specifies its own network stack from the physical layer to the application layer. The main applications of Z-Wave include home automation such as garage doors and alarm systems, among other things. The application layers have been tailored to suit those needs, hence the technology is quite specific to that market segment. Connecting Z-Wave devices to the Internet has to be done through a protocol translation gateway. [10]

6LoWPAN [6] is a competing standard to ZigBee that has the added benefit of interoperability with other IP-based systems. 6LoWPAN introduces an adaptation layer between the link and network layers, allowing IPv6 packets to be transmitted over IEEE 802.15.4 based networks. Moreover, it defines a header encoding to support fragmentation and compression in case the datagrams do not fit within a single frame. [3]

3.1 Bluetooth Low Energy

The standard Bluetooth radio has been widely implemented and is available in mobile phones, laptop computers, audio headsets and many other devices today. Bluetooth Low Energy (BLE) is a low energy variant of the standard that enables the use of the interface with resource-constrained devices such as sensors [4]. The low power variant of Bluetooth was introduced in revision 4.0 of the Bluetooth specifications, was enhanced in Bluetooth 4.1, and has been developed even further in successive versions. IPv6 over Bluetooth LE is dependent on Bluetooth 4.1 or newer. BLE is designed for transferring small amounts of data infrequently at modest data rates at a very low cost per bit, which makes it attractive especially for Internet of Things applications. Bluetooth LE has a possible range of over 100 meters.

Every Bluetooth LE device is identified by a 48-bit device address. The link layer uses star topology, consisting of a central router and peripheral nodes. The router can connect to multiple peripherals, and is assumed be less constrained than the peripherals. Direct communication only takes place

between a central and a peripheral, thus there is no direct communication between peripherals. In a primary deployment scenario, the central router will act both as a 6LoWPAN border router and a 6LoWPAN node.

Bluetooth LE nodes have to find each other and establish a link-layer connection before any IP-layer communications can take place. The Bluetooth LE technology sets strict requirements for low power consumption and thus limits the allowed protocol overhead. 6LoWPAN standards provide useful functionality for reducing overhead which can be applied to Bluetooth LE. This functionality comprises of link-local IPv6 addresses and stateless IPv6 address autoconfiguration, neighbor discovery and header compression.

A significant difference between IEEE 802.15.4 and Bluetooth LE is that the former supports both star and mesh topology and requires a routing protocol, whereas Bluetooth LE does not currently support the formation of multi-hop networks at the link layer.

Bluetooth Special Interest Group has introduced two trademarks to be used: Bluetooth Smart refers to single-mode devices that only support Bluetooth LE, whereas Bluetooth Smart Ready devices support both Bluetooth and Bluetooth LE. [4]

3.2 Near Field Communication

Near Field Communication (NFC) is a set of standards for portable devices to establish radio communication with each other by bringing them to close proximity. NFC typically requires a distance of 10 cm or less between the devices. Devices using NFC may be active or passive: passive ones only able to send information, while active devices can send and receive data. The NFC technology has been widely implemented and is available especially in mobile phones and laptop computers. NFC always involves an initiator and a target, the communication is node-to-node only. The initiator generates a radio frequency field that can power a passive target, which enables passive NFC targets to be applied on very small form factors such as stickers or key cards that do not require any batteries. NFC also supports bidirectional communication, provided that both peers are active.

One of the differences between IEEE 802.15.4 and NFC is that the former supports both star and mesh topology, whereas NFC can only support direct peer-to-peer connection and simple mesh-like topology because of very short transmission distance. Due to this characteristic, 6LoWPAN functionality, such as addressing, auto-configuration and header compression, is specialized into NFC. [2]

4 Comparison

4.1 Energy efficiency

Energy in sensor networks is provided either by battery or by scavanging energy from the environment, such as solar power: in either case, it's a constrained resource. For radio-equipped sensors, the radio transceiver is the most power-consuming component as Figure 4 shows. It also shows that there is little difference in consumption between receiving and transmitting mode operation. Therefore it is undesirable

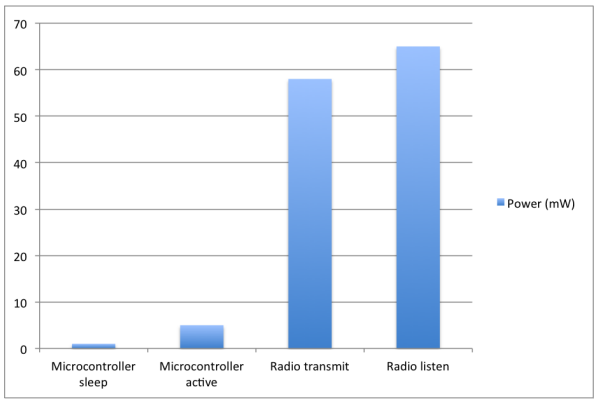

Figure 4: Power consumption of the microcontroller and radio transceiver on a Tmote Sky prototyping board. [10]

to have the radio always on due to power constraints, and consequently different approaches to radio resource management have been investigated.

The IEEE 802.15.4 standard supports many features that result in significant power savings. However, achieving a desired data rate and maximizing the lifetime of individual sensors are often conflicting goals. One of the most important energy efficiency features is the possibility of turning the transceivers off most of the time and activating them only when required; idle listening is a major source of energy waste. Depending on the device, it can be in various sleep mode states that have a different impact on the energy consumption and the time it takes for the device to wake up [10]. For applications with timing constraints, timely delivery may be more crucial than energy saving.

The energy efficiency of the ZigBee standard is mainly at the physical and MAC layers. ZigBee supports two operating modes: one that is very effective but limited in scope to star topology, the other which basically tries to reduce power consumption by using very low duty cycles. ZigBee's low power consumption limits transmission distances to 10-100 meters line-of-sight, depending on power output and environmental characteristics. [9, 1]

4.2 Security

Smart objects typically have slightly different security- and threat models than general purpose computing systems due to the their various applications. Security is important for smart objects because they are often deployed in important infrastructures such as the electrical power grid. They are often deployed in places that make them vulnerable to intrusion attempts and in places where security breaches can not be tolerated. Sensors used in home automation can lead to intrusion attempts. It's also been shown that remote reprogramming of pacemakers has been possible, making security deficiences possibly even lethal. [10] A widely accepted model for determining security consists of confidentiality, integrity and availability.

Confidentiality is perhaps the most evident notion of security: data is confidential only and only if the right parties

can access it. Confindentiality is not easy to ensure for sensor networks as data is transmitted mainly via wireless channels. Authentication in low-power networks is challenging as the system is non-centralized and there is no central server to verify identities [10]. Security models for low-power devices in general require a way to securely transmit keys to the sensors. As a solution to this, a model called the resurrecting duckling model [8] has been developed, where devices are imprinted to match their mothers, which they then blindly follow.

Integrity is kept if it the data can not be tampered with before reaching its recipient. Even though integrity and confidentiality are related to each other, they are different concepts. Availability makes sure that the data is available to the right parties at the time it is needed. If availability is breached, the system is said to be suffering from a Denial of Service (DoS) attack. For low-power devices, availability can be breached by so-called sleep deprevation attacks: an attacker may spoof the device into keeping its radio on and thus depleting its battery. [10]

Security is often confused with encryption. Even though encryption is an important part of security, strong encryption alone is not a security model - most breaches happen due to other problems than cryptographic failures [10]. Applications of LoWPAN often require confidentiality and integrity protection, which present their own challenges. For example, the microcontrollers used in low-power sensors may not be able to execute asymmetric decryption operations within a reasonable time, hence the encryption algorithms have to be computationally efficient. IEEE 802.15.4 mandates link-layer security based on AES, but omits high-level details such as key management. Both ZigBee and 6LoWPAN use 128-bit AES encryption.

ZigBee includes methods for key establishment and exchange, frame protection and device management. ZigBee's AES encryption is possible at network or device level. Network level encryption is achieved by using a common network key; device level encryption by using unique link keys between pairs of devices. End-to-end security is provided so that the keys are shared only between the source and destination nodes, routing can be applied independent of security considerations. [9, 1]

The transmission of IPv6 over Bluetooth LE has similar requirements for security as for IEEE 802.15.4. Bluetooth LE Link Layer supports encryption and authentication by using the Counter with CBC-MAC mechanism [11] and 128-bit AES block cipher. Upper layer security mechanisms may exploit this functionality when it is available. Key management in Bluetooth LE is provided by the Security Manager Protocol. [4]

NFC's short communication distance can be considered as a security advantage as it makes third-party eavesdropping difficult - any attempt to hack into the RF between the NFC devices must happen within the 10 cm operating radius [2].

4.3 Interoperability

Interoperability is the ability of systems from different vendors to operate together. It is one of the leading factors when choosing a wireless protocol; applications should not need to know the constraints of the physical links that carry their packets. Especially for sensor networks it is essential as the devices emerge at a large scale. At the physical layer, sensors must agree on matters such as the frequencies of transmission, type of modulation and the data rate. At the network level, nodes must agree on the data format and the addressing of nodes, as well as how the network topology functions.

ZigBee builds on the 802.15.4 standard and defines new upper layers on top of it. This means ZigBee devices can interoperate with other ZigBee devices, assuming they utilize the same profile. [9, 1] 6LoWPAN offers interoperability with other wireless 802.15.4 devices as well as with devices on any other IP network link with a simple bridge device. Bridging between ZigBee and non-ZigBee networks requires a more complex application layer gateway.

5 Conclusion

The Internet of Things is still emerging with applications ranging from home automation to medical and even to military applications - making the field extremely broad and the number of use cases and operating conditions wide-ranging. Prior to the consensus of adopting IP for low-power devices came a reality, several non-IP based solutions were deployed, such as ZigBee and Z-Wave. ZigBee gained significant popularity and has been widely adopted by the industry, as it became the most mature solution.

There are a number of implementations for IPv6 over low-power networks, but not all of them directly compete with each other and most are still being actively developed. Adopting IPv6 to low-power networks has several challenges that these technologies have tried to overcome; the low-power requirement makes it important for the underlying network to be fault-tolerant and easily configurable, while still supporting a large number of connected nodes. 6LoWPAN, the effort to adapt IPv6 for low-power networks, is attractive for its interoperability and it has been used by various technologies such as Bluetooth Low Energy. Unlike 6LoWPAN, ZigBee cannot easily communicate with other protocols. NFC is ideal for extremely short-range and secure communication and the possibility to have completely passive nodes offers numerous use cases.

References

[1] P. Baronti, P. Pillai, V. W. Chook, S. Chessa, A. Gotta, and Y. F. Hu. Wireless sensor networks: A survey on the state of the art and the 802.15.4 and ZigBee standards. *Computer communications*, 30(7):1655–1695, 2007.

[2] Y.-G. Hong, Y.-H. Choi, J.-S. Youn, D.-K. Kim, and J.-H. Choi. Transmission of IPv6 Packets over Near Field Communication. *Active Internet Draft*, 2015.

[3] J. Hui and P. Thubert. Compression format for IPv6 datagrams over IEEE 802.15. 4-based networks. 2011.

[4] M. Isomaki, J. Nieminen, C. Gomez, Z. Shelby, and T. Savolainen. Transmission of IPv6 Packets over BLUETOOTH Low Energy. 2015.

[5] N. Kushalnagar, G. Montenegro, C. Schumacher, et al. IPv6 over low-power wireless personal area networks (6LoWPANs): overview, assumptions, problem statement, and goals. *RFC4919, August*, 10, 2007.

[6] G. Montenegro, N. Kushalnagar, J. Hui, and D. Culler. Transmission of IPv6 packets over IEEE 802.15.4 networks. *Internet proposed standard RFC*, 4944, 2007.

[7] M. Siekkinen, M. Hiienkari, J. K. Nurminen, and J. Nieminen. How low energy is bluetooth low energy? comparative measurements with zigbee/802.15. 4. In *Wireless Communications and Networking Conference Workshops (WCNCW), 2012 IEEE*, pages 232–237. IEEE, 2012.

[8] F. Stajano. The resurrecting duckling. In *Security Protocols*, pages 183–194. Springer, 2000.

[9] The ZigBee Alliance. Zigbee specification, 2006.

[10] J.-P. Vasseur and A. Dunkels. *Interconnecting Smart Objects with IP: The Next Internet*. Morgan Kaufmann, 2010.

[11] D. Whiting, N. Ferguson, and R. Housley. Counter with CBC-MAC (CCM). 2003.

New applications to reduce energy consumption of cellular network using Smart Grid

Toni Mustajärvi
Student number: 84026K
toni.mustajarvi@aalto.fi

Abstract

In this paper energy consumption of mobile cellular base stations is studied with an aim to find new ways to reduce network power consumption. Base stations uses a lot of energy and power costs for operators is significant. Additionally usage of cellular network is rapidly increasing and thus forcing the operators to deploy new base stations. This paper tries to find new ways the whole cellular network could save energy especially using the new smart grid without going deeply to network technology. Discussed solutions are delayed content distribution, communication with content providers about how the data can be transmitted, automatic shutdown of base stations and payments based on the used power type.

KEYWORDS: base station, cellular network, smart grid, energy saving methods

1 Introduction

Mobile cellular base stations (BS) are the most energy utilizing components in the cellular network [7]. Have been shown that more than 80% of the energy consumption in operators cellular network power usage originates from the network. From this more than 50% of comes from the base stations [2]. Also when taking in to account the fact that usage and bandwidth requirement in a mobile cellular network is rising [2], there is a need for more energy efficient solutions.

Smart grids are rapidly replacing traditional power grids that are powering cellular networks and base stations. These grids offer new possibilities for a network to adapt when trying to lower the costs and keeping power usage as optimal as possible. For example, smart grids can have a varying energy price at hourly level depending on the network usage. Also they can predict the costs in the future which can be used by the cellular network to create usage models. Smart grid also provides information about energy producers and their emission levels which is vital information if green energy is the important aspect.

In this paper new theoretical applications are discussed to improve the cellular network energy efficiency. Focus is for solutions that can be achieved programmability with help of the smart grid. First the paper discusses about the smart grid and its characteristics and how the base stations have been using energy and what it means if they include own power sources. Finally four new solutions are discussed how the energy consumption of the base stations could be improved.

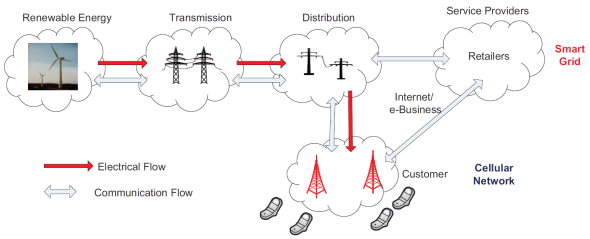

Figure 1: Diagram of cellular network with smart grid [2]

2 Smart Grid and Cellular Network Energy Consumpion

In this chapter new smart grids are discussed and new possibilities rising from them is evaluated against cellular networks.

2.1 Smart Grid

Smart grid is the next generation of the power grid system. It incorporates information flow in the grid and by doing so adds many new possibilities for operating and using the grid. Electricity consumption has been increasing but additionally demand for renewable energy has also increased. For example, with smart grid it is possible to select the electricity provider instantly or apply dynamic changing pricing of the electricity based on grid total consumption and other variables. Same also applies to energy producers who can adapt to changing energy requirements by lowering or increasing the energy generation. Smart grid provides realtime information about the grid which can be used by all the connected parties. [2]

Smart grid also brings new interesting opportunity for normal energy consumers to sell energy in to the grid. For example if consumer is able to store energy and grid is in need of energy, consumer can sell their stored energy to the grid.

Figure 1 presents a simplified format of the smart grid used by the cellular network. Usual power flow is from the grid to cellular network but there are situation when this is also reversed if the network has battery power available and

grid is in a need for extra energy. In the figure red white arrows represent the information flow smart grid support and provides. This information is provider to service providers but also for the consumers.

2.2 Smart Grid and varying energy price

Power grids are heavily shifting from traditional network grids towards smart grids. Smart grid brings more intelligence to power grids and allows new possibilities to optimize power usage. The smart grid controls the power generation and distribution based on the demand. They also allow users to choose the energy source, for example base stations could buy renewable energy produces by wind turbines.

One of the main characteristic of smart grids is the varying energy price. In traditional power grids, prices were determined a long beforehand. Because of this, prices consumers saw were very steady and no sudden shifts were possible. Smart grids on the other hand have the information flow incorporated to the grid and this allows rapid changes in the energy price and methods how the power is generated. It is even possible to have a negative energy price when there is much more energy produced than consumed [4]. This is possible because it is not always feasible to shut down power plants which can be very expensive operation and bringing the plant on again could take a long time. This dynamic nature of smart grids can be exploited in cellular networks. One possible use case for this is discussed in a section 3.1.

2.3 Base station energy generation and storage

Base station are not only tied to power grid for providing necessary energy for it to operate. Base station can have its own power source such as solar panels or wind turbines. This also means the base station needs to be able to store the energy to tolerate changing power requirements and energy generation fluctuations. These technologies allows new ways to utilize smart grid when managing the base stations [6]. This aspect is dicussed more in the section 3.4.

3 Solutions for energy-efficient cellular networks

In this chapter different promising solutions for utilizing smart-grids to achieve better energy consumption regarding the base station usage. Content is mostly in theoretical level.

3.1 Delayed content distribution

It is known that data transfer rates in the mobile cellular network are growing rapidly [3]. There are many applications to affect this but mostly this is due to the rapid growth of smart phone markets. New smart phones are better and better of consuming huge data and video streams.

A common theme for these applications is that they require the data instantly because user is waiting the content to be delivered. However there are many tasks that can be done in the background and the scheduling the time is not so strict. One good example of this kind of data is mobile phone software updates. Currently the phone operating system updates can take almost 1GB of data and transferring this amount information to many phones in the cellular network takes huge amount of bandwidth. Games and other applications are also growing in size and updates frequently happen. Considering mobile devices the most energy consuming operation is the active network connection. To limit the network activity over longer periods, content can also be loaded as fast as possible into memory and consuming can happen in much slower speed. Clear advantage of this is the possibility of closing the connection thus saving a lot of battery. [8]

With the new intelligent smart grid, these data transfers could be scheduled to happen when it is the most suitable for the network. There are many things that must be taken into consideration when scheduling the data transfers with the aim to reduce the total energy consumption of the cellular network.

First the grid needs to know how and when the transfer should be commenced at the latest. To achieve this, distributors needs to provide this information for the smart grid. This aspect is discussed more in the section 3.2.

Second the grid should take into account the current network usage. The cellular base stations have a characteristic where the power consumption does not rise linearly with the used capacity [1]. Instead they consume almost full power even when utilized lightly. Because of this it would be efficient to use the remaining capacity to something useful, for example transmitting the updates and other not time sensitive information. On the other hand, the quality of the network should be guaranteed to other users where data access is needed immediately.

Third the energy price and emission levels from its generation should be taken into account. Because the smart grid can predict the energy price in the future it is possible to schedule data transfers to time where the price is the lowest. For example if the price is low during nights where also the usage is in the lowest points, instead of shutting down the cellular nodes, the low energy price could be used to transfer the updates. Of course this automated process should be studied more and it should have a good algorithm to determine the optimum usage of the cellular network. Another side is that base stations can themselves choose the energy provider taking into account the type of production and prices. Especially the emission level of the energy producer can be one of the important criterion and thus should be taken into account when choosing the provider [2].

When green and renewable energy sources are prioritized, they should also be added in the calculations.

Although previous discussion mostly concentrated on the mobile phone usage, the same principles and ideas can be applied to other use cases. One good example would be the automated data collection devices in the network. They could be instructed by the network to send the information in the specific time of the day.

3.2 Communication with content providers

Cellular network and base stations cannot by themselves determine the type of the content and how it can be deliv-

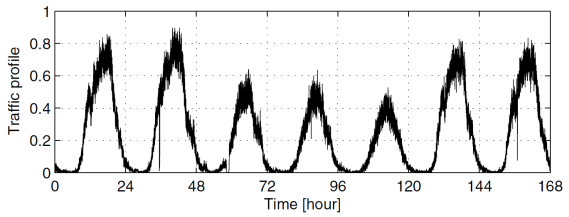

Figure 2: Cellular network normalized traffic profile [1]

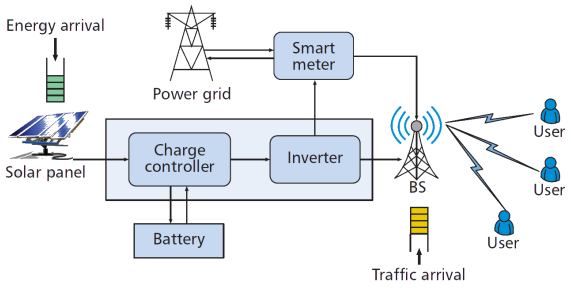

Figure 3: Battery powered base station [5]

ered. They need this metadata information from the content providers.

This information is even more critical to have when considering the real-time nature of the smart grids. Situation can be constantly changing and for example when the grid generates more energy than it is consumed at the time base stations could utilize this information to do more work. But for the network to know what work can and should be saved for later use, it needs information about this. One simplified idea could be to mark the content to support a different delivery tactic with information such as, should be delivered within a week. Then with this information, the cellular networks and individual base stations could delay the sending in more optimum time.

3.3 Automatic shut down of base stations

Because base station energy utilization is not linear and even low utilization takes almost the full power usage, shutting down base stations during low traffic times is a good possibility to save energy.

Another important note is the traffic profile of the cellular network during the week. An example profile is presented in the figure 2. Traffic is at the highest rate in the middle of the day when people are working. During nights, the traffic is at the lowest point. Also weekends have slightly lower traffic compared with working days. Density of deployed base stations needs to be high in more populated areas and centers where people are working. Considering the fact that base stations power consumption correlates very little with the used capacity, there is a lot of potential to shutdown some of the base stations during low traffic times. This however is not possible to achieve when considering only single base station but instead the whole cellular network grid needs to be taken into consideration. Communication between base stations is needed.

Shutting down base stations is not only done because traffic is low but it can be done also when energy prices in the smart grid are too high or the pollutant level associated with the provider are considered too high [2]. The smart grid enables the cellular network to adapt in these fast changing situations by providing necessary information about the grid and energy being produced [5].

When base stations are shut down in the network, network needs to also consider the coverage of the network and quality of service for its users. Coordinated multipoint (CoMP) is a technology to make sure the network requirements are satisfied [2].

3.4 Battery power assisted base stations

One of the main benefits of the smart grid is the support for the two-way flow of the energy. This means that the energy generated or stored in the base station can also be sold to the grid if there is demand for it. Combined with the real-time information about the prices and providers it is possible to create battery-powered base stations that can themselves determine the optimum combination of energy usage from the grid and the base station's own energy reserve.

In figure 3 battery powered base station is presented. Power grid provides the constant flow of energy to the base station. Additionally solar panel is used to generate energy on the site and it is stored in the battery. Using intelligent power management algorithms the base station can determine which energy to use.

Cellular network usage fluctuates during the day and there can be large peaks in the usage. Normally, energy for these is drawn from the energy grid but a battery supported base station can use its own energy reserves. This gives possibility to level the energy consumption and with it the price of the energy. The same situation can happen in the smart grid. The price can suddenly increase for a short time, for example because of sudden over demand for energy than its produced. During these situations, the base station can also rely on its own battery capacity. Additionally it is also possible to sell the available energy to the grid. Leithon [6] showed that significant cost savings are possible with this technology. There are however many parameters to consider, for example the battery capacity, charging rates and the correlation between price and consumption profiles.

The base station can also be equipped with its own energy source like solar panels. This has been proved to be a working solution in Mobile World Congress 2010 where 100% solar power base station was deployed [9]. Using intelligent energy management algorithms the solar panels are able to maintain the battery power during the base stations operation.

3.5 Payments based on used power type

The mart grid allows very precise information about the energy providers, how the energy is generated and pollution levels for the provider. This information could be provided also for the end users of the network. It would allow new

types of contract for the cellular operators to provide network access where energy is generated with green energy. The same concept has already been applied to consumers and households by the traditional energy providers.

For example offer two similar contract in technical perspective but the difference would be is the energy produced with renewable methods. This might affect the prices if renewable energy is costlier which usually is the case but many people are more energy friendlier and are willing to pay more for green solutions.

4 Discussion

On this paper, we have discussed possible new ways to have more energy efficient cellular networks and base stations with the help of the new smart grid.

Bandwith requirement in the mobile cellular networks is rapidly increasing [3]. This is mostly due to the fact that smart phones and other mobile devices are better and better of consuming the right media. This produces more demand on the cellular network and how the energy consumption can be optimized with the increasing traffic.

Internet of things (IOT) is a coming trend in the markets and basic idea behind it is that every device is connected to the internet. This allows the new possibilities of control and automation in many areas. These devices however needs to have wireless communication methods of some kind. Cellular networks are the most suitable candidate for most use cases. These devices most probably can cause a large amount of traffic in the cellular network making it considerably harder to achieve energy savings if traffic keeps increasing. This aspect needs to be taken into account when designing the new green base stations. Cellular network and individual base stations should have some sort of control over the IOT devices and how they communicate with the internet. Always open connections drastically lowers the possibilities of achieving energy savings in the base stations.

5 Limitations

This paper was written by conducting a literature review and collecting possible solutions in the section 3 to improve the energy efficiency of cellular network base stations powered by the smart grid. A limited number of articles were found around this topic and it might affect the results. Also some of the solutions were not based on any specific study but ideas were drawn from similar topics.

6 Conclusion

In this paper cellular networks powered with new the smart grid was discussed. We found out that the new smart grid open new possibilities for cellular networks to adapt their work to achieve energy savings and lower emissions. Smart grids offer near the real-time information flow of the energy prices and consumption. With this information even individual base stations can select which operator to buy the electricity.

Delayed content distribution with the help of content providers could provide large energy savings when they are working together. Larger file transfers done in the mobile cellular networks could be scheduled for time when the energy price is low or base stations have plenty of green energy stored in its own batteries.

On the other hand base stations could optimize their own operation with the help of the smart grid and intelligent power management algorithms. Utilizing energy generation and batteries, base stations have been demonstrated to be able to work completely on their own. Also base stations could be operated in network level to allow shutting down of the base station in places where utilization is low and work could be divided into other nearby stations.

Finally the pricing based on the green energy used in the network could be provided for the customers. With smart grids base stations are able to select the providers based on numerous parameters. Providing new pricing options would be possible.

Future studies should be concluded about how the cellular network could adapt for the varying energy prices when powered with smart grids. For this study very limited material was found and none of the papers described this aspect. Secondly, the delayed content distribution should be studied and how it could be utilized with the smart grid where energy prices are changing heavily during the day.

References

[1] O. Blume, H. Eckhardt, S. Klein, E. Kuehn, and W. M. Wajda. Energy savings in mobile networks based on adaptation to traffic statistics. *Bell Labs Technical Journal*, 15(2):77–94, Aug. 2010.

[2] S. Bu, F. R. Yu, Y. Cai, and X. P. Liu. When the smart grid meets energy-efficient communications: Green wireless cellular networks powered by the smart grid. *IEEE Transactions on Wireless Communications*, pages 1–11, 2012.

[3] Y. Chen, S. Zhang, S. Xu, and G. Li. Fundamental tradeoffs on green wireless networks. *IEEE Communications Magazine*, 49(6):30–37, June 2011.

[4] F. Genoese, M. Genoese, and M. Wietschel. Occurrence of negative prices on the German spot market for electricity and their influence on balancing power markets. pages 1–6. IEEE, June 2010.

[5] T. Han and N. Ansari. Powering mobile networks with green energy. *IEEE Wireless Communications*, 21(1):90–96, Feb. 2014.

[6] J. Leithon, Sumei Sun, and Teng Joon Lim. Energy management strategies for base stations powered by the smart grid. In *Global Communications Conference (GLOBECOM), 2013 IEEE*, pages 2635–2640. IEEE, Dec. 2013.

[7] N. Nupponen. Energy efficiency development of lte network authentication protocol. Master's thesis, Aalto

University - School of Electrical Engineering, September 2013.

[8] K. Pentikousis. In search of energy-efficient mobile networking. *IEEE Communications Magazine*, 48(1):95–103, Jan. 2010.

[9] D. Valerdi, Q. Zhu, K. Exadaktylos, S. Xia, M. Arranz, R. Liu, and D. Xu. Intelligent energy managed service for green base stations. In *GLOBECOM Workshops (GC Wkshps), 2010 IEEE*, pages 1453–1457. IEEE, Dec. 2010.

Security and privacy in smart energy communities

Kari Niiranen
Student number: 66925J
kari.j.niiranen@aalto.fi

Abstract

Smart energy has been rolled out largely with a focus on possible energy savings and efficiency improvements. However privacy and security aspect has not usually been considered during this deployment. This paper looks at the current situation and highlights possible privacy and security concerns. Also proposed solutions are looked at and the future of privacy and security in smart energy grids and communities is discussed.

KEYWORDS: security, privacy, smart energy

1 Introduction

Smart energy grids aim to solve problems traditional energy distribution experience due to the increased demand and the introduction of renewable sources. Traditional grids react slowly to changing power requirements and can lead to situations where parts of the grid have an abundance of energy while others are starved of it. Smart energy grids solve these problems by permitting the grid to react to the changing situations faster and automatically[1]. Smart energy grids can provide an option for customers who generate their own energy to sell any excess back to the grid, while allowing the customer to buy energy to cover peak requirements that the customers own generation capacity can not handle.

In addition to the smart grid, modern home appliances can form a network of smart devices that are able to communicate with each other and with the grid. This means that appliances with high energy requirements can ask the grid about current energy costs and inform the grid of future energy requirements. Additionally, the grid can inform the customers appliances about possible usage limitations leading to a situation where the networked appliances decide which of them has priority over others. This can help to lower the utility bill by only running these appliances when the energy is cheap. Also it is able to improve the efficiency of the grid and can help in situations where the grid would otherwise become overloaded and fail. By combining many neighbouring networks, a smart energy community is formed where many customers can work together to save energy. This could allow the community to purchase shared generating capacity, thus further reducing their utility bill.

As smart energy communities and usage of smart home appliances become more widespread, security and privacy issues are going to become more widespread too. This is because most migrations from the traditional grid to a smart grid are done to reduce costs and improve efficiency, security and privacy are usually an afterthought for the utility company. For example, a malicious individual might use monitoring tools provided by the utility to find out when a home or business is empty. Insecure appliances and smart meters might allow an attacker to lie about energy consumption and production, or change the power state of an appliance. However, when used correctly the same tools allow energy saving, more efficient distribution of power, and could help save money on the electricity bill[8].

The main challenge is therefore providing the users with accurate information about their energy consumption while at the same time keeping the same information secure from others. Additionally, the same information should be hard to modify or at least modifications should be obvious to prevent fraud.

This paper looks at proposed and current solutions to privacy and security issues smart grids experience. Section 2 gives an overview of currently used solutions and their security risks. Section 3 focuses on proposed solutions and finally section 4 discusses the future of smart energy.

2 Current solutions in practice and security risks

In this section we are going to look at what current solutions are in use and what security risks they face. The smart grids provide many improvements to how energy usage is controlled and monitored. However this also means that there are new ways to abuse the system and violate the customers privacy. Additionally the smart grid allows the customer to change role from a pure consumer to a consumer-producer that can provide additional energy generation capacity by connecting their own solar panels or other source of energy to the grid.

2.1 Tracking of energy usage

Accurate metering of energy consumption is important as this information is used to bill the customer. In traditional grids, the meter was read by a technician at certain intervals. This meant that the customer was either billed based on estimated energy consumption or the customer sent the meter reading to the utility company. If the amount the technician read from the meter differed from the estimated or reported consumption, the customer was either billed for any extra usage if the estimate was below actual usage or given credit if the estimate was over actual consumption. Modern smart grids use remote readable smart meters that store accurate usage information, with detailed consumption statistics

available with a precision of at least one hour. This has the advantage that the utility company can offer to bill the customer by hourly electricity price, thus possibly reducing the utility bill. Even if the customer is billed a fixed sum per each kilowatt hour used, the customer is billed every time for exactly the amount of energy used. This means there is no need to adjust the billing in case of incorrect estimates.

Some utility companies, such as Helen (previously known as Helsingin Energia) and Fortum in Finland, even provide the customers an interface where it is possible to see and track accurate usage information. An example screen shot from the Helen's Sävel plus service is shown in figure 1. This allows the customer to see how new appliances affect energy usage in near real time. However, this same information can be used to build a profile of the customers habits. For example lower energy consumption could be interpreted as the usage location being empty.

By providing accurate usage information to the customer, the customer may use this information to compete with others in how much energy they are able to save. For example Opower[10] provide methods for customers of an utility company to be more engaged with their consumption of energy. While they recently concluded that a stand alone application connected to the utility company's data did not provide the widespread results they had hoped, they believe that behavioural change and customer engagement is the way to reduce energy utilisation. For larger customers who might manage multiple buildings, there exists similar solutions. For example, Granlund provides a software solution called Granlund Manager[6] that performs energy consumption monitoring on the building level.

Additionally this all relies on the remote readable meter to be secure. For this to be true, automated testing frameworks should be available. Dantas et al. propose in their paper [4] a tester that is based on sending random data in the meters communications channel and seeing how the meter and connected systems cope with incorrect commands. This allows hardening the system against deliberate attempts to crash monitoring systems by feeding them unexpected data. Also it can reveal possible vulnerabilities that could allow unauthorised alteration of usage data.

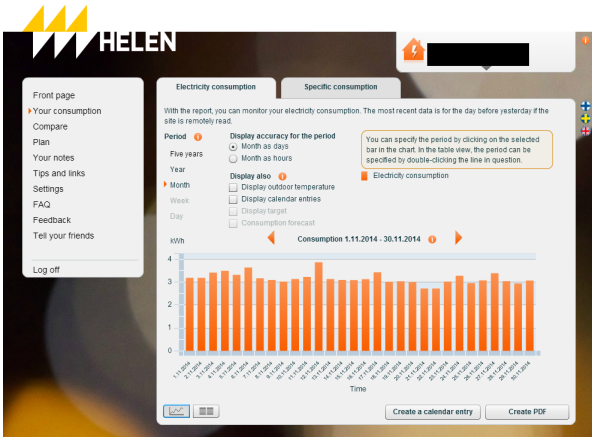

Figure 1: Screenshot from Helen's Sävel plus service.

2.2 Smart appliances

Smart appliances allow remote control or automation of their functions by the user. For example the user could instruct the washer to start washing clothes when the solar panels reach a certain energy generation level or if that never happens, at a predetermined time. Thus the user has clean laundry waiting when returning home. Other possibilities include sequencing of different appliances to keep energy usage below a certain threshold. This is especially important when generating own energy and trying to avoid buying energy from the grid if at all possible. Also appliances could communicate with the grid and only run when energy is cheap or has enough spare capacity.

However, considering how current networked smart devices are often vulnerable to exploits and can expose user information to the Internet, this has possible security and privacy implications. For example remote controllable appliances should use authentication that is unique to that specific unit and not one shared with all units of the same model. Or even worse, there is no authentication at all. Second issue is that these kind of items tend to be configured once and then forgotten, thus resulting in vulnerable versions of the firmware being used long after a fixed version has been made available. Usually the user is completely oblivious to the fact that the cheap webcam he installed to watch the pet during work is actually publicly available on the Internet[7] or that his new smart television is used to spy on him[9]. Additionally, the user may connect their air conditioning to their smart network, thus allowing an attacker to cause physical discomfort by turning the temperature to extreme values during the users sleep.

Thus new appliances should be made easy to configure with automatic updates of the firmware enabled by default. However, this poses the risk of the new firmware accidentally breaking the appliance. This does not remove the issue for old, no longer supported appliances as these won't receive new updates fixing security issues. And even though the appliance may no longer be supported, it may still work for its intended purpose, thus making it questionable why it should be replaced just because the software it is running is no longer maintained.

Other risks smart appliances can pose is that even if the management interface is secure, the appliance might be installed in some permanent manner. Thus if the current owner sells the property, the new owners might not realise that the appliance can be managed remotely. This issue has no easy solutions as the user usually wants appliances to do their function without any additional hassle. Thus requiring the user to change password at a regular interval can lead to dissatisfied customers for the manufacturer. If there is a way of resetting the password, it needs to be hidden enough that it is not accidentally triggered, which leads to the issue that the new owners most likely are not aware of the existence of such a function.

2.3 Smart energy communities

Smart energy communities are communities where a group of nearby people decide to work together for a shared goal. A smart energy community might decide to lower the over-

all energy bill by purchasing a solar panel and selling any excess generated power to the grid. Any profit generated is then put into purchasing more solar panels or other improvements to shared areas, such as a new grill for the building's yard. As the aim is to lower the whole community's energy usage, usually some sort of usage tracking is provided. This could be an adaptation of the monitoring system detailed in section 2.1. To encourage lowering energy usage, a ranking system could be employed where one could track their energy savings compared to others.

However, ranking systems have privacy concerns. If everyone's information is available publicly, then it could be used to see if somebody is not at home. Also, if the target is to save energy as a community, rankings could cause conflicts if someone does not meet the targets set. The simple solution would be to only show the user their own rank and numbers. To make this more secure, the data used to build this ranking should be handled in a manner that preserves the privacy of individual members of the community. One way of doing this could be to use Sharemind[3] or a similar system. The aim is that after running the data through the system to generate the rankings, it is hard to derive from the rankings which member of the community that data corresponds to.

3 Proposed research solutions

3.1 Obfuscation of information

To solve the issue of reporting exact usage data to the smart energy community, obfuscated data is sent. The obfuscation is implemented using an algorithm that guarantees that the actual numbers can not be easily extracted from the sent value, but when used in calculations the correct totals can be derived. Dimitriou and Karame propose [5] a method of obfuscating amounts and prices different consumer-producers in a smart energy grid are willing to buy and sell energy at. This is done so that members of the grid are unable to easily game the system for monetary gains.

The same system could be utilised to protect the numbers in the smart energy community detailed in section 2.3. This would move the storage of actual values to the user and storing only the obfuscated values centrally for calculations. Doing it this way protects the system from accidental information leakage. Other way of doing this is adapting the research done by Bogdanov et al. where they used Sharemind to combine two different data sources in a way that preserved the privacy of the individuals in said data sources[2].

The problem with these solutions are that they cost money. Thus the cost of adopting either or both methods needs to be low enough that everyone can adopt them without the cost being the deciding factor. Additionally the average person might not see the benefits of either system. The problem with proposing such a solution is that it can be seen as stating that other involved parties can not be trusted. Thus the argument for implementing a solution like these needs to be construed carefully.

3.2 Standardisation of appliances

As there are almost as many different management interfaces for smart appliances as there are different manufacturers, a user might not want to learn every appliances own special interface. By providing standard ways of managing certain common functions, such as resetting user credentials or enabling automatic firmware updates, the user is more likely to actually perform said action. Thus it is more likely that the user keeps the appliance more secure. However this does not guarantee that the user actually performs the required actions to keep the appliance secure.

The standardisation of interfaces is important also when buying a property or used appliances as it is likely the manuals for the appliance are lost. If the interface is standardised, the appliance can still be managed without difficulty assuming the user has sometime used an appliance with the standardised interface.

The problem with this approach is that manufacturers are not keen on sharing functionality with others.

4 The future

In the future, everything will be connected to a network. This means that unless care is taken to ensure information is stored securely, user privacy and security could easily be compromised. Looking back at earlier technologies, security and privacy have been secondary concerns that only get dealt with when they are exploited. Going forward, it is very likely that the first implementations are almost guaranteed to contain privacy and security issues that are solved, when the system becomes more mature. However, people tend to avoid making the same mistakes twice, thus new appliances do not usually contain the same issues as older models. For example, wireless access points no longer use a default unencrypted wireless network like the devices from early 2000 did. Modern access points use encrypted networks by default to protect users who tend to plug devices in and forget about them.

When working with privacy and security, one has to maintain a fine balance between security and usability. In most cases by making a system more secure, the usability suffers and vice versa. Thus the question arises of how much damage to the involved parties can this information cause if it leaks instead of how secure the system is. Thus requiring multiple authentication sources to remotely adjust the thermostat of the air conditioning would be prohibitive to the usability of that function. Even if someone managed to access that without rights, the worst that happens is that the inhabitants feel uncomfortable for a while and some excess energy might be used.

References

[1] S. M. Amin and B. F. Wollenberg. Toward a smart grid: power delivery for the 21st century. *Power and Energy Magazine, IEEE*, 3(5):34–41, 2005.

[2] D. Bogdanov, L. Kamm, S. Laur, P. Pruulmann-Vengerfeldt, R. Talviste, and J. Willemson. Privacy-

preserving statistical data analysis on federated databases. In *Proceedings of the Annual Privacy Forum. APF'14*, volume 8450 of *LNCS*, pages 30–55. Springer, 2014.

[3] Cybernetica AS. Sharemind. `https://sharemind.cyber.ee/`, 2015. [Online; accessed 17-April-2015].

[4] H. Dantas, Z. Erkin, C. Doerr, R. Hallie, and G. v. d. Bij. efuzz: A fuzzer for dlms/cosem electricity meters. In *Proceedings of the 2Nd Workshop on Smart Energy Grid Security*, SEGS '14, pages 31–38, New York, NY, USA, 2014. ACM.

[5] T. Dimitriou and G. Karame. Privacy-friendly planning of energy distribution in smart grids. In *Proceedings of the 2Nd Workshop on Smart Energy Grid Security*, SEGS '14, pages 1–6, New York, NY, USA, 2014. ACM.

[6] Granlund. Granlund manager. `http://www.granlundmanager.fi/`, 2015. [Online; accessed 11-April-2015].

[7] C. Heffner. Exploiting surveillance cameras - like a hollywood hacker. In *Black Hat USA 2013*, February 2013.

[8] P. McDaniel and S. McLaughlin. Security and privacy challenges in the smart grid. *Security & Privacy, IEEE*, 7(3):75–77, 2009.

[9] B. Michéle and A. Karpow. Watch and be watched: Compromising all smart tv generations. In *Consumer Communications and Networking Conference (CCNC), 2014 IEEE 11th*, pages 351–356. IEEE, 2014.

[10] Opower. Opower. `http://opower.com/`, 2015. [Online; accessed 10-April-2015].

Software technology challenges in 3D printing

Ari Oinonen
Student number: 53768V
aoinonen@gmail.com

Abstract

3D printing is predicted to revolutionise manufacturing by enabling custom-made products at low cost. This seminar paper explores the usage of software technology related to 3D printing. The various steps in 3D printing, from acquiring a 3D model to the actual fabrication, are covered from a software technology perspective. The biggest challenges in 3D printing technology might be on the hardware side, but software technology will also play a big role in the advancements of the technology. Combining individual research and knowledge will be necessary in order to achieve a highly advanced printing process.

KEYWORDS: software technology, software technology challenges, 3d printing, additive manufacturing

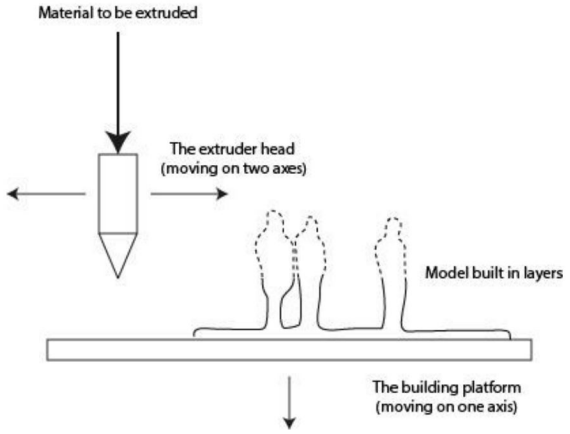

Figure 1: The operating principle of Fused Deposition Modelling. [20]

1 Introduction

1.1 Objective of the paper

The objective of this seminar paper is to look into 3D printing from a software technology perspective. The main topics covered include the various steps in 3D printing where software technology is used, and the main software challenges in those areas.

1.2 Overview of 3D printing

3D printing is often called additive manufacturing (AM), referring to the fabrication method in which the object is built from scratch. This is as opposed to substractive manufacturing methods with CNC (Computer numerical control) machines such as milling, where the process starts from a block of material that is then partly removed until only the desired shape remains. [20]

3D printing is predicted to revolutionise manufacturing by enabling custom-made products at low cost [16]. In 2012, the global market for AM products and services increased by 28.6 % from the previous year, reaching USD $2.2 billion [26].

1.3 3D printers

3D printing technology has been around since the 1980s. However, it was not until the 2000s that 3D printers arrived on the entry level market. Affordable, desktop-sized, personal 3D printers allow access to additive manufacturing capabilities for those of all skill and interest levels, including amateurs and hobbyists along with those focused on professional usage. Research about 3D printers has primarily focused on professional machines, although more recent research covers entry level 3D printers as well. [9, 20]

3D printing can be done using several different techniques. Two common ones are fused deposition modelling (FDM) and selective laser sintering (SLS). Fused deposition modelling printers have an extruder head that moves on two axes, while the building platform moves on one axis. The operating principle can be seen in Fig. 1. A heated thermoplastic material is extruded to build the object in layers. Selective laser sintering also manufactures the object in layers, but in a different fashion - the pre-determined shape is formed by using a laser beam to cure a photo-cureable resin powder in a container, and afterwards the unused powder surrounding the cured object can then be collected and used in another printing. [20]

Various materials can be used to 3D print objects. Entry level machines generally use ABS (acrylonitrile butadiene styrene) or PLA (polylactic acid) plastics as their material. Some commercial printers are also able to print objects using metals [16].

As synthetic materials have an environmental impact, environmentally friendly printing with wood flour and wood pulp is being tested [10]. Ceramics and edible materials are also being printed [16]. A great deal of research is being done in the medical field, but according to Bose et al., we are still a long way from completely printing functioning hu-

man tissue [3].

3D printing can be subcontracted on the Internet or from brick-and-mortar businesses, such as MR MAKE in Germany. However, subcontracting is slower and the costs are higher. [20]

2 3D printing process

2.1 Overview

The 3D printing process is made up of three main phases: 3D model creation, process planning and the actual fabrication of the physical part or parts.

he 3D model can be created using a CAD (Computer-aided design) tool, generated with scripts or scanned in 3D from a real-life object. The model may have errors or any number of issues that prevent it from being 3D printed or the printing might not be optimal. The model geometry can be corrected, it can be balanced and orientated for optimal printing performance. Section 2.2 looks at the various methods and approaches of creating 3D models for printing.

The model may have errors or any number of issues that prevent it from being properly 3D printed, causing the printing to be sub-optimal, or simply not function as planned. The model geometry should be corrected if need be, as it can be balanced and orientated for optimal printing performance.

The printing process must be thoroughly planned before the actual printing. This includes slicing the model and generating tool-paths for the printer. This results in machine instructions called G-code that are given to the printer. The printing process is covered in section 2.3.

The actual fabrication of the physical part is done based on the G-code instructions. There are various ways the object can be fabricated.

The model usually needs manual post-processing after it is printed, where support structures and any defects on the surface have to be removed. The model might also be painted.

2.2 3D model creation

3D printing starts by creating a 3D model for which there are two main approaches. One approach is to create the 3D model from scratch by using a CAD (Computer-aided design) tool or by generating the model with code. The second approach is to create the model from images or point clouds based on a real-life physical object.

The model you need may have already been created by someone else and been made available for download on a website such as Thingiverse. Thingiverse is a website for sharing of 3D models that are meant to be created physically using 3D printers, laser cutters, milling machines and many other technologies. [25]

2.2.1 Computer-aided design tools

A traditional method for creating 3D models is using a CAD tool. Common CAD tools have not been designed to support special requirements that 3D printing and specific 3D printers have. For example, certain surface angles and shapes may not print well on some 3D printers.

```
/* [Customize body] */

//Set the outside length of your pencil box.
length=190;//[70:400]
//Set the outside depth of your pencil box.
depth=70;//[50:400]
//Set the total height of your pencil box.  The top of the box is set at 15mm.
//Extra height is added to the body section.
height=40;//[40:150]

//Choose divider orientation.  Long is for the X direction.
long = 1;//[0,1,2]
//Short is for the Y direction.
short = 2;//[0,1,2,3]
//When you have 2 long dividers,
// picking yes here will put short dividers in the center section.
center = 0;//[1:Yes,0:No]
 .
 .
 .
module body(fn=20) {
    difference(){
        union(){
            case(z-15,fn);
            hinge(0,z-15,180,fn);
            translate([0,-y/2+5.75+wall,(z-10)/2+t])
                cube([13,12,z-20-2*t],center=true);
            intersection(){
                translate([0,-y/2+5.75+wall,5])
                    cube([13,12,10],center=true);
                translate([0,-y/2+rout+wall,rout/1.5])
                    rotate([0,90,0])
                        cylinder(r=rout,h=13,center=true,$fn=fn);
            }
            if (long>0){
                translate([0,-(long-1)*yin/6,0]){
                    for (i=[1:long])
                        translate([0,(i-1)*yin/3,0])
                            divider(xin,z-15,2,fn);
                }
                if (short>0){
                    translate([-(short-1)*xin/6,-yin/(long+1)*(long/2),0])
                        for (i=[1:short])
                            translate([(i-1)*xin/3,0,0])
                                rotate([0,0,90])
                                    divider(yin/(long+1),z-15,1,fn);
                }
            }
```

Figure 2: .scad script: excerpt of global parameters ① and a module ② [1]

2.2.2 Scripting and customisable models

Most CAD tools are not designed to support parameters in models. However, customisable models can be created by writing scripts. One of 3D printing's biggest strengths is to be able to offer personalised designs with minimal waste, as well as its customisability offering the possibility of altering the model to fit the individual needs of people.

OpenSCAD is a non-visual, programmer-oriented solid modelling tool. It is much like a 3D compiler, rendering 3D models from script files. [24] Thingiverse extended the OpenSCAD language to allow users to design parametric objects by adding various parameters in the script. The end user can then change the parameters, such as dimensions or optional parts, customising the object to themselves. An example of the OpenSCAD language and parameters can be seen in Fig. 2. [1]

Acher et al. conducted a study of Thingiverse in order to establish a possible connection with software product line (SPL) engineering. They found hints that SPL-alike techniques are used in 3D printing. However, there are many limitations in OpenSCAD parameters, with OpenSCAD not having a mechanism for specifying constraints between parameters or restricting some values of a parameter. SPL techniques can improve the reusability of 3D models, but the

complexity induced by the tools and languages can reduce the intended benefits. Reusing parts of a model and creating reusable building blocks for new models may turn out to be a powerful method for additive manufacturing, but the possibilities appear rather limited for the time being. [1]

2.2.3 3D scanning

3D scanning can be done using photographs or videos of an object. Medical imaging data can also be used. [15] [18]

Photogrammetry is the technology of deriving 3D data, characteristics, and attributes from 2D images [20]. Madracevic et al. present a method that focuses on detecting, grouping, and extracting features, such as edges and faces, present in a given picture and trying to interpret them as 3D clues [15]. Scanning the whole surface of a 3D object can be done using three methods. One method is to rotate the object on a rotating stage while a single camera is stationary, a second method is to move the camera around a still object, and a third method is to use multiple cameras to scan the object from multiple angles simultaniously. [11]

Another approach to creating a 3D model is to gather point clouds with a laser scanner and create a model from the data. Static terrestial laser scanning (TLS) can produce dense point clouds and can reach an accuracy of 8 mm or better in ranges of 10 m to 50 m. Mobile laser scanning is an emerging technique, in which a laser scanning system is mounted on a vehicle. Point clouds obtained using laser scanning techniques are typically very large in size, and can contain hundreds of millions of points. Processing them requires a lot of computational resources and can be time consuming. Typical problems encountered with this approach include gaps in the data, often caused by obstructions in the scanning, and varying point densities in the data set. [20]

2.2.4 Error checking

Especially when scanned, the model may need some correcting and adjusting before it can be used. It may have surfaces that do not connect, or undesirable gaps in the model. The mechanics of the 3D printer or method used can also cause some complications on the printing of the model. For example, walls or shapes at some specific angles can potentially have issues printing optimally due to the physical limitations of the mechanical process. [20]

2.2.5 Balancing and hollowing

Balancing or hollowing the model or object can also be necessary, as the object may not stand on its own when printed, it may need too much material to be practical, or be too heavy. [13] [6]

The cost of materials used for 3D printing remains high, despite 3D printers becoming popular even for home users, and solutions are desirable. Wang et al. present an automatic solution for designing a skin-frame structure for 3D objects, with the resulting object consisting of a hollow frame inside the object, covered with a solid skin on the outside. An example of this can be seen in Fig. 3. The frame structure generated by the algorithm is guaranteed to be physically stable and printable. The biggest limitation for their solution

Figure 3: Printed model with part of the skin removed to show the inside structure. [22]

appears to be scaling the solution for large objects that exceed the tray size of the printer. In this case, segmentation of the objects would be required. Assembling parts of various frame structures while maintaining strength and stiffness may turn out to be challenging in some situations. [22]

2.3 Process planning

Before the physical fabrication process, the geometric model undergoes a process planning, in which it is converted into instructions for the printer. In general, this means first converting the model to an STL file format, optimising the position of the model, slicing the model into layers, and finally calculating the tool-path for the printer, which is transferred to the printer as G-code instructions. [7]

2.3.1 File formats

The physical fabrication of the model is done by layering continuous slices on top of each other. The Standard Tesselation Language (STL) file format is a standard format used by AM machines as an input to generate the slices. Some of the reasons for the popularity are the simplicity of the format and ease of file generation without requiring complicated CAD software [4]. The STL format uses planar triangular facets to approximate the surfaces of the part, which introduces errors in the part representation, especially in highly curved surfaces. This approximation leads to errors in parts manufactured by AM machines. The approximation can be reduced by increasing the tesselation and using smaller triangles, which increases the number of triangles exponentially, leading to increased file size and possibly slowing down the entire process. [17]

ASTM, an international standards organisation, recently introduced a new file format, Additive Manufacturing File (AMF). It utilises curved triangles based on second degree Hermitive curves. The curved triangles are sub-divided back to planar triangles for slicing and the same approximation error might still happen. Anand et al. introduces a new file format based on Steiner patches, that are used also in the slic-

ing stage. Steiner patches are bounded Roman surfaces that can be parametrically represented by rational Bezier equations. Steiner surfaces are of higher order and thus the file format is more accurate. Their results show that the Steiner format is able to significantly reduce geometric dimensioning and tolerancing (GD&T) errors in parts manufactured by AM processes compared to the STL and AMF formats. The new format is currently incomplete, with many aspects still requiring further research and development. [17]

2.3.2 Build orientation and placement

Build orientation, or the orientation that the object will be printed in, is one of the most important process planning tasks, since it directly affects many factors such as surface quality, build time, cost, and the complexity of the required support structures. Since the printing is done in layers and each layer is generally produced in an optimal tool-path pattern, the surface qualities of the object might vary depending on the orientation. [28, 19]

Armilotta et al. developed a method for the optimal selection of build orientation based on widely accepted selection criteria. Since the build orientation is an important factor, there are many things to consider when choosing the optimal orientation. Their method consosts of a two-step selection procedure, where near-optimal orientations are generated and evaluated visually. The method was implemented in an interactive tool. Compared to existing methods, their solution focuses on improving the surface quality. Further work is still needed to take into account the different aspects in which surface quality can be specified, as well as improving balancing between the various selection criteria. [2]

Zhang et al. studied the orientation optimisation of multi-part production, where a group of parts in the same build chamber should be simultaneously oriented optimally. Their solution first generates a set of finite optimal alternative orientations for each part, then a genetic algorithm is applied to search for an optimal combination of orientations to minimise the total build time and cost at a global optimal level. Further research is yet needed for many aspects of the problem. [28]

Dedoussis et al. developed algorithms for simultaneous fabrication of multiple parts, where the layout of the parts on the printer platform require optimisation. The solution assumes the build orientation to be already fixed. They employed a Genetic Algorithm technique for the 2D nesting of parts on the platform. Their solution leads to satisfactory layouts, leading to substantial improvement of 3D printer utilisation, leading to time and cost savings. [5]

2.3.3 Support structures

Printers based on Fused Deposition Modelling (FDM) technology require support structures in order to print the models, with support structures connecting overhanging parts with lower parts of the object or ground. Since the support material needs to be printed first and then discarded, optimising its volume leads to savings in material and printing time. Vanek et al. introduce a geometry-based optimisation framework for reducing the need for support structures. An example of the support structures can be seen in Fig. 4. [19]

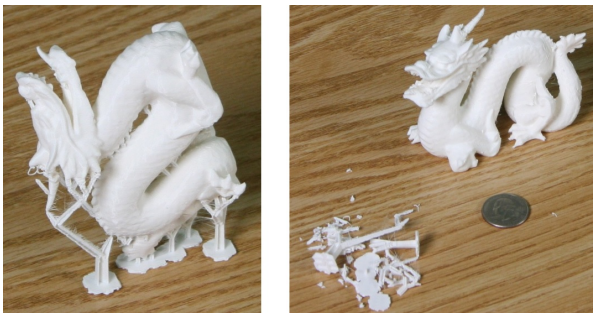

Figure 4: Left: Support structures generated by Vanek et al. Right: The finished object after the support structures have been removed. [19]

Orientation is a major factor here, causing more support to be needed if the object is not oriented optimally. Their approach starts by orienting the object into a position where only a minimal area requires support. Next, tree- like support structures are generated for the points that require support, while attempting to minimise the overall length of the support structures. Vanek et al. manage to reduce the printing time and the amount of material by over 10 % compared to previously existing solutions. [19]

According to them, the most serious limitation of the approach is that it does not provide any structural evaluation, being purely geometry-based. Also, it is developed for just one printer model, so some parameters would probably have to be varied for a more general solution. [19]

2.3.4 Slicing the model

The physical fabrication of the model in 3D printing is done in overlapping flat layers. The data for the object comes from a 3D model, usually represented by a triangle mesh. The geometric data is divided into layers in a step called slicing, with many different strategies for slicing meshes existing. According to Gregori et al., most of the current literature is concerned with issues such as the quality of the model, specific improvements in the slicing process and memory usage. Gregori et al. approach the problem from an algorithmic complexity perspective, proposing an algorithm that is asymptotically optimal under certain common assumptions. [7]

The manufacturing speed of 3D printers is still very slow, often limiting usage. To achieve the best quality result, an object is sliced uniformly with the finest resolution of the printer. Wang et al. present an adaptive slicing method where the thickness of the layers varies while the printed result preserves the visual quality of the print with the finest resolution. The solution might however cause minor visual artifacts along the vertical boundaries. [21]

2.3.5 Generating the tool-path and G-code instructions

3D printers cannot actually print an object without a special algorithm that creates the computer numerical control (CNC) instructions, commonly known as G-code. Brown et

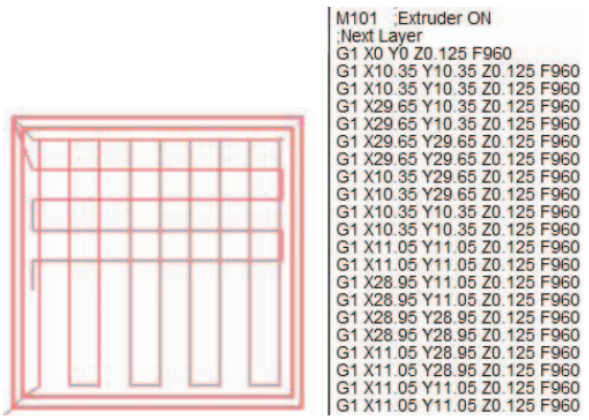

Figure 5: Sample G-code and its representation. [4]

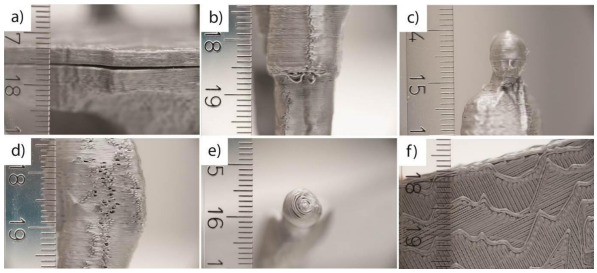

Figure 6: Various artifacts caused by 3D printers: (a) Seam; (b) Hanging layers; (c) Visible layers; (d) Gaps in layers; (e) Over/under extrusion; (f) Step pattern. [20]

al. present an algorithm for slicing and tool-path generation, resulting in a G-code file for an entry-level 3D printer. The work is part of a larger effort focusing on the development of a low cost AM platform. [4]

Tool-paths are generated from the layers. Brown et al. use a crisscross based approach to generating the tool-paths. The crisscross pattern can be made with various slope angles. The type of fill and tool-path, such as the slope angle used, can affect the fabrication time and quality. [4]

G-codes are commands that control the 3D printer and each G-code has a distinct function. Instructions such as the printers' extruder speed and temperature also have to be considered. An example of G-code can be seen in Fig. 5. [4]

2.4 Printing the model

There are two common methods for the actual physical fabrication of the model. Fused deposition modelling produces the model by extruding small beads of material which harden immediately to form layers. Another approach is selective fusing of materials in a granular bed. [23]

The resulting object might have various defects, such as uneven surfaces. [20] These will be covered in more detail in section 2.5. The typical 3D printing process is done in an open-loop manner, in which the printing process is determined in advance and does not change during the process. Lu Lu et al. seeked to fix this problem in ink-jet 3D printers by developing a closed-loop layer-to-layer control algorithm, measuring the height of the printed layers and adjusting the printing accordingly. Their experimental results show that using the closed- loop algorithm improves the printing quality, resulting in more consistent shapes and smoother surfaces. [14]

Lee et al. developed a hybrid rapid prototyping system combining 3D printing and machining in a five-axis machine tool. The approach makes it possible to work on the object after 3D printing in order to achieve more accurate dimensions or better surface finish. In addition, it removes the need for support structures and improves the build time by around 50 % in the tests performed. [12] This combined approach opens up a lot of new possibilities with the combination of added axis and machining capabilities. Getting the most out of these possibilities will require new approaches for software as well.

2.5 Post-processing

The model usually needs post-processing after it is printed. Possible support structures and any defects on the surface must be removed. The model may also be painted, since the 3D printing is often done using only a single colour. Examples of artifacts caused by 3D printers can be seen in Fig. 6. [20]

2.6 Multiple colours

Several printers have multiple extruders, which allows objects to be formed from multiple materials or colours. The extruders are mounted side by side on the printer. However, according to Hergel et al., the print quality suffers when objects with color patters are printed. The most severe issue is the oozing of plastic from the idle extruders, causing plastics of different colours to bleed onto each other. There are multiple ways to improve the quality, but they do not come without downsides. Hergel et al. aim to solve the issue with software, introducing three techniques that complement each other in improving the print quality significantly. They first reduce the impact of oozing plastic by choosing a better orientation for the part. Secondly, they build a disposable rampart in close proximity of the part, giving the extruders an opportunity to wipe oozing strings. Finally, they make a tool-path generation algorithm that avoids and hides most of the defects due to oozing and seamlessly integrating the rampart. According to them, the work makes it easier to get satisfactory results with multiple color prints. [8]

Zhang et al. developed a 3D true color printing robot, including a specialised file format and network communication protocol. They successfully printed a 3D true colour model of the terrain of Taiwan. [27]

3 Conclusion

3D printing technology has been around since the 1980s. However, it was not until the 2000s that the 3D printers arrived on the entry level market. The 3D printing process

consists of multiple detailed steps, many of which require algorithms and optimisation problem solving. These steps include, but are not limited to 3D scanning, customisation of the model, orienting and positioning the model, preparing the model for printing and converting it to instructions for the 3D printer.

The biggest challenges in 3D printing technology may be on the hardware side, but software technology will also play a big role in the advancements of the technology. Even the basic steps of the process have not yet been implemented perfectly. Most research done appears to focus on a single point of the process. However, the steps are connected and will often require changes in other parts of the process as well. Combining all this research and knowledge will be necessary in order to achieve a highly advanced printing process.

3D printing applications are demanded in many fields which all have their specialised requirements. Printing in the medical field is likely the most demanding of them all, with a great deal of specialised software needed before we are able to conquer the greatest challenges, such as printing fully functioning human organs.

References

[1] M. Acher, B. Baudry, O. Barais, and J. . Jézéquel. Customization and 3d printing: A challenging playground for software product lines. *ACM International Conference Proceeding Series*, 1:142–146, 2014.

[2] A. Armillotta, M. Cavallaro, and S. Minnella. A tool for computer-aided orientation selection in additive manufacturing processes. pages 469–475, 2014.

[3] S. Bose, S. Vahabzadeh, and A. Bandyopadhyay. Bone tissue engineering using 3d printing. *Materials Today*, 16(12):496 – 504, 2013.

[4] A. Brown and D. De Beer. Development of a stereolithography (stl) slicing and g-code generation algorithm for an entry level 3-d printer. 2013.

[5] V. Canellidis, J. Giannatsis, and V. Dedoussis. Efficient parts nesting schemes for improving stereolithography utilization. *CAD Computer Aided Design*, 45(5):875–886, 2013.

[6] A. N. Christiansen, R. Schmidt, and J. A. Bærentzen. Automatic balancing of 3d models. *CAD Computer Aided Design*, 58:236–241, 2015.

[7] R. Gregori, N. Volpato, R. Minetto, and M. Silva. Slicing triangle meshes: An asymptotically optimal algorithm. pages 252–255, 2014.

[8] J. Hergel and S. Lefebvre. Clean color: Improving multi-filament 3d prints. *Computer Graphics Forum*, 33(2):469–478, 2014.

[9] D. P. Industry. History of 3d printing: The free beginner's guide, 2014. [Online; accessed 7-February-2015].

[10] L. R. Julien Gardan. 3d printing device for numerical control machine and wood deposition. *Int. Journal of Engineering Research and Applications*, 4:123–131, dec 2014.

[11] S. Klein, M. Avery, G. Adams, S. Pollard, and S. Simske. From scan to print: 3d printing as a means for replication. *HP Laboratories Technical Report*, (30), 2014.

[12] W.-C. Lee, C.-C. Wei, and S.-C. Chung. Development of a hybrid rapid prototyping system using low-cost fused deposition modeling and five-axis machining. *Journal of Materials Processing Technology*, 214(11):2366–2374, 2014.

[13] L. Lu, A. Sharf, H. Zhao, Y. Wei, Q. Fan, X. Chen, Y. Savoye, C. Tu, D. Cohen-Or, and B. Chen. Build-to-last: Strength to weight 3d printed objects. *ACM Transactions on Computer Systems*, 33(4), 2014.

[14] L. Lu, J. Zheng, and S. Mishra. A layer-to-layer model and feedback control of ink-jet 3-d printing. *IEEE/ASME Transactions on Mechatronics*, 2014. Article in Press.

[15] L. Madračević and S. Šogorić. 3d modeling from 2d images. pages 1351–1356, 2010.

[16] N. Notman. High precision 3d printing of metals warms up. *Materials Today*, 18(1):5 – 6, 2015.

[17] R. Paul and S. Anand. A new steiner patch based file format for additive manufacturing processes. *CAD Computer Aided Design*, 63:86–100, 2015.

[18] G. Taubin, D. Moreno, and D. Lanman. 3d scanning for personal 3d printing: Build your own desktop 3d scanner. 2014.

[19] J. Vanek, J. Galicia, and B. Benes. Clever support: Efficient support structure generation for digital fabrication. *Computer Graphics Forum*, 33(5):117–125, 2014.

[20] J.-P. Virtanen, H. Hyyppä, M. Kurkela, M. Vaaja, P. Alho, and J. Hyyppä. Rapid prototyping - a tool for presenting 3-dimensional digital models produced by terrestrial laser scanning. *ISPRS International Journal of Geo-Information*, 3(3):871–890, 2014.

[21] W. Wang, H. Chao, J. Tong, Z. Yang, X. Tong, H. Li, X. Liu, and L. Liu. Saliency-preserving slicing optimization for effective 3d printing. *Computer Graphics Forum*, 2015. Article in Press.

[22] W. Wang, T. Wang, Z. Yang, L. Liu, X. Tong, W. Tong, J. Deng, F. Chen, and X. Liu. Cost-effective printing of 3d objects with skin-frame structures. *ACM Transactions on Graphics*, 32(6), 2013.

[23] Wikipedia. 3d printing — wikipedia, the free encyclopedia, 2015. [Online; accessed 7-February-2015].

[24] Wikipedia. Openscad — wikipedia, the free encyclopedia, 2015. [Online; accessed 20-March-2015].

[25] Wikipedia. Thingiverse — wikipedia, the free encyclopedia, 2015. [Online; accessed 19-March-2015].

[26] C.-C. Yeh. Trend analysis for the market and application development of 3d printing. *International Journal of Automation and Smart Technology*, 4(1):1–3, 2014.

[27] S. Zhang, H. Wang, X. Chen, H. Qian, J. Liu, and S. Lin. Design and implementation of modular software system for three-dimensional true color printing robot. pages 6454–6459, 2013.

[28] Y. Zhang, A. Bernard, R. Harik, and K. Karunakaran. Build orientation optimization for multi-part production in additive manufacturing. *Journal of Intelligent Manufacturing*, 2015. Article in Press.

MOOCs and Authentication

Jan Pennekamp
Student number: 472 706
Jan.Pennekamp@aalto.fi

Abstract

With the increasing number of Massive Open Online Courses (MOOCs) and corresponding students, the value of MOOCs for professional careers is gaining greater attention. Currently, there are some MOOC providers that already offer certificates by using advanced authentication approaches. However, these certificates are usually not accepted officially, because the students are not authenticated properly.

This paper examines existing and proposed approaches for MOOC authentication. Furthermore, it evaluates their feasibility and usefulness from the poinf of view of the students, the providers and the authorities. Based on this evaluation, the paper identifies five key aspects that are important for proper MOOC authentication: Trade-off between Security and usability, Privacy Concerns, Inaccuracies, Time of Authentication, and Costs.
KEYWORDS: MOOC, authentication, feasibility, signature track, SSO, stylometry, proctoring

1 Introduction

The first time Massive Open Online Courses (MOOCs) appeared, was in 2008. After the unexpected success (2,000 enrolled students) of the first MOOC offered by the University of Manitoba, Canada, more Canadian institutions started experimenting with this new approach to online learning. The first approaches basically offered their regular university courses to non-affiliated students in form of video recordings and online exercise submissions. The U.S. based universities followed this trend in 2011. The first U.S. MOOC attracted over 2,500 students (University of Illinois Springfield). From that point on, the first MOOC providers, such as Khan Academy, were founded. They also offer their own content that is not directly related to university courses. Due to the rapid interest, more providers followed in 2012: for example, Coursera, edX, and Udacity. Section 2.1 gives a more detailed introduction and categorization. To highlight the growth, Coursera offers a good impression: In 2013, there were already over 3.7 million registered students [16]. The content, the target audience, and the required level of expertise of the offered MOOCs differs between different providers and courses.

The main reason for students to enroll in a MOOC is obtaining knowledge that they can later reference in their professional career. In developed countries, education is the key qualification for getting a job. At the same time the costs for education increases in these countries [16]. Hence, students are exploring new approaches, such as MOOCs, to improve their CV. Furthermore, there are fundamental differences when comparing MOOCs with regular university courses. The students can work through the lectures in their own speed at their preferred location at their preferred time. Thus, they can arrange their studies according to their convenience. The reason why companies and universities do not accept MOOCs officially, is that, in theory, anyone could have taken the course. Although some MOOC providers offer enhanced authentication and identification services, the courses are not accepted yet as replacements of university courses in a curriculum. Hence, the following questions arise: Do we really need authentication and who benefits from it?

There are several reasons for introducing authentication. First, the existing approach of providing certificates after passing a course would be properly validated. The offered certificates would acquire an official status. Therefore, MOOCs could be used in official capacities: e.g., as selection tests or as an accepted form of higher education [17]. Second, proper authentication utilizes the benefits of MOOCs even more: large parts of the population have access to the Internet and therefore have access to "free" higher education. Naturally, providers might introduce fees to cover the costs of implementation and operation of their offered authentication service. However, while this is a downside to the openness of MOOCs, the gathered data about participating students might be improved in contrast to non-authenticated MOOCs, because more serious customers participate in the courses and at the same time cheating is contained more effectively. Finally, increasing the amount of gathered valuable personal information could compensate the costs. Nowadays, utilizing personal information is part of the business model of various successful companies such as Facebook, Google, and Twitter.

2 Background

This section introduces MOOCs in general and well-known MOOC providers. Furthermore, it defines what tamper-proof authentication is and what challenges have to be tackled to ensure this. The following sections build on the presented information.

2.1 MOOCs

Massive Open Online Courses (MOOCs) are courses of study that are accessed via the Internet and are usually free of charge. Hence, they target a high number of participants. Additional fee-based services, such as certification or tutor-

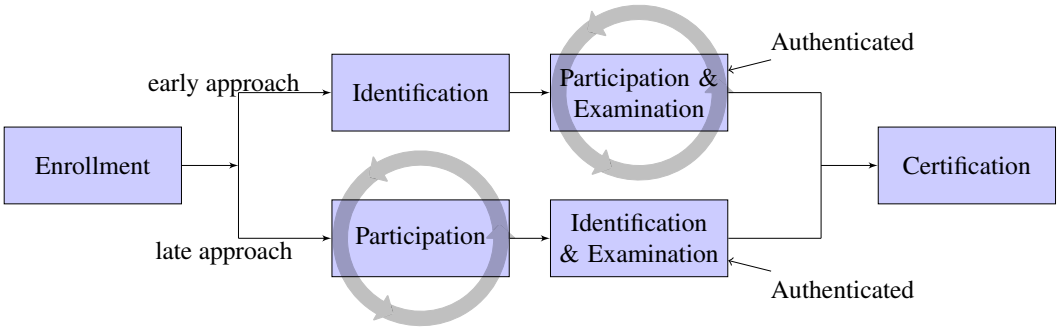

Figure 1: Usual Authentication Process used by Presented MOOC Providers

ing services, are offered as well. There are both commercial and non-profit providers. A large number of the offered courses are based on regular university courses that have been adapted (in terms of assessment) to operate as a massive open online course. Most course providers reward the participant with an certificate of participation upon completion. MOOCs are expected to have an enormous impact on higher education in the future [16].

In recent years, various MOOC providers have established themselves in the market [5]. The main characteristic of MOOC providers is their commercial focus, but also the license under which they share course content is used to set apart. The major commercial MOOC provider are Coursera[1], Udacity[2] and Udemy[3]. EdX[4] and FutureLearn[5]. They are examples of non-commercial providers that are affiliated with various well-known universities. Besides these, there are providers which offer their content under a free license, for instance Khan Academy[6] and P2PU[7]. In Finland, both University of Helsinki and Aalto University offer courses on MOOC platforms as well. They use `http://mooc.fi` and Eliademy[8].

2.2 Authentication

Authentication can be divided into two subtasks [10]. First, identity verification, which is used to link a profile a student created at a MOOC provider to a real-world identity. This verification usually has to take place once, because afterwards, the provider stores this matching. A common approach is to check the submitted information against an ID card or student records. This only results in successful verification if the records are available to the provider. Second, identity authentication is used to check that the verified student is actually the student participating in the MOOC. Ideally, this authentication should take place whenever the student submits information that is used for grading. This check is necessary to ensure that the enrolled student is fulfilling the course's tasks. Currently, identity authentication is mostly performed based on rather weak criteria, such as password authentication. This mechanism does not prevent cheating at all. However, there are situations (e.g. exams) for which the providers require a stronger authentication. In the following, we will refer to both subtasks as authentication.

A common sequence of events in MOOCs is illustrated in Figure 1. As discussed later, there are currently two approaches when providing authentication in MOOCs. First, the early approach offers effectively authenticated participation and examination. The identity verification occurs already after course enrollment, hence, the student is properly authenticated course-long. Second, the late approach only authenticates the examination securely. The course submissions before the examination are not linked with advanced authentication, hence, only a fraction of the student's commitment is authenticated properly. Finally, both approaches end with grading. Certificates are handed out whenever offered by the provider.

The stakeholders that are interested in authentication can be divided into three major groups. A first group are the students who have an interest in ensuring that their performance is valued and accepted in their subsequent professional career. For that reason, many students will even accept privacy intrusive mechanisms for authentication that will be presented in Section 3. A second group of stakeholders are companies and universities that want to be able to officially acknowledge the effort some applicants invested in the past. With the rising number of people wanting to pursue an academic education, MOOCs would be an alternative to traditional selection tests. As MOOCs offer academic education to a new part of the population, they could be very useful if the involved parties would credit the invested effort. A third group are the MOOC providers that are interested in providing proper authentication to their customers, because more students ultimately mean more value creation. In addition, the commercial providers have the chance to charge their customers for these authentication services. The value of MOOCs to the customer mostly correlates with the acceptance of the certificates by companies and universities.

[1] `https://www.coursera.org/`
[2] `https://www.udacity.com/`
[3] `https://www.udemy.com/`
[4] `https://www.edx.org/`
[5] `https://www.futurelearn.com/`
[6] `https://www.khanacademy.org/`
[7] `https://p2pu.org/`
[8] `https://eliademy.com/`

2.3 Challenge

The general challenge is to determine how tamper-proof the authentication and identification offered by MOOC providers for their customers has to be. This raises the question at what point a MOOC is no longer a MOOC course, because, for example, the person has to be physically present at the MOOC provider to verify his identity. At which point does the student stop using MOOCs: Be it due to privacy issues or authentication overhead? Will it be acceptable for the stakeholders that there are still ways to obtain a certificate with illegal help/cheating, while accepting provided certificates as valid proof of accomplishment?

Section 3 considers some answers to these questions by presenting both existing and proposed approaches. Whenever possible, the paper refers to major MOOC providers that were listed in Section 2.1.

3 MOOC Authentication

Currently, there is no solution to the challenge of proper authentication for MOOCs, because the existing authentication approaches detract from some of the advantages introduced by MOOCs: Either they are not accessible by all interested students or they are not exclusively on the Internet or they are not free. The remaining approaches are not judged as insufficiently trustworthy. Otherwise, the cases of officially accepted MOOCs would be significantly higher. The proposed solutions either are not usable in their current state or they are too costly to deploy. Nevertheless, we will present these approaches and discuss the corresponding advantages and drawbacks.

3.1 Existing Approaches

The most important existing approaches the providers use for authenticating students at MOOCs are presented in this section. Additionally, the identified drawbacks of each approach are presented.

Proctored Exams is the most traditional approach when comparing MOOCs to regular university courses. In order to participate in a proctored exam, the student has to enroll for the exam and then participate in an exam which is offered on-site at the MOOC provider or in one of the regional testing centers. Before taking the exam, the student's identity is verified in the same way as it is done in traditional university exams [14]. Providers which offer this approach are edX and Udacity.

This approach seems to be quite "secure" in ensuring that the right student takes the exam. However, there are certainly some disadvantages as well. First, these proctored exams are only offered at a limited number of locations at certain times: This possibility is usually limited to the United States. Furthermore, this approach is costly for the student: Udacity charges 89 $ per exam [14] and edX 95 $ [1]. Hence, the openness and online factor of MOOCs are negated. Students from around the world do not have the money to participate in these proctored exams (travel, distance). For an already employed student, the fixed time is an is issue as well. Overall, the student's flexibility is limited. However, also the MOOC provider has expenses, because proctors have to be available to invigilate the exams. Finally, this approach is not completely secure as there is no way to verify that the student participating in the exam is the same student who participated in the MOOC. This might be an issue for a traditional university course as well.

(Cousera) Signature Track is another approach which was introduced by Cousera in 2013. Signature Track uses the student's typing pattern to identify him. A special phrase is used to create the typing pattern. This should uniquely identify the student in later sessions. The typing pattern is linked to the student's account as soon as the student verifies its identity with an ID. This identity verification is performed using a webcam. Whenever the student wants to use the platform, he has to authenticate by typing the special phrase and take a current picture with a webcam [7].

This approach is used to authenticate the student throughout the course. However, there is no protection against the contribution of "third-parties", such as fellow students or the Internet, throughout the course or in examinations [1]. As it is not possible to verify the typing patterns individually, this step is performed by software. Additionally, the picture verification of whether the same student is currently participating in the course is done automatically as well. This process reveals multiple drawbacks. First, it has not been extensively evaluated how accurate or inaccurate the typing and photo matching is [10]: Can another student imitate the typing pattern? Krause [8] states that he was able to develop an implementation which is robust against imitation attempts. Is a photograph in front of the webcam sufficient for authentication? Do varying lightning conditions affect authentication? Second, not every student participating in MOOCs has access to a webcam whenever studying. Third, Signature track costs 30 to 100 $ per course [1]. These points clearly limit the course's reach. Finally, students might not accept such a privacy-invasive approach.

Single-Sign-On (SSO) is another approach to offer authentication to the student participating in a MOOC. In theory, this approach works well, because the student does not have to register for each provider, but rather can use an existing account (identity) to authenticate [15]. However, in practice, there is no global single-sign-on solution that offers proper authentication. Hence, only students known to and already identified at the provider can log in. As most MOOCs were offered by universities, their students could use a SSO system to authenticate: Their university credentials are used for granting access. A related method, which is open to the public is OpenID: A decentralized authentication system that is used to provide user information to co-operating providers. P2PU offers login authentication based on OpenID.

Major MOOC providers do not offer a SSO based on university credentials anymore. Eliademy, for example, grants access to users having a Microsoft Live account. EdX implements this approach for Google and Facebook accounts. While this approach does not restrict the target audience significantly, it does not provide proper authentication, because the identity is not really verified. The SSO approach basi-

cally limits the access to MOOCs to a certain group that is already identified. Currently, this is usually the case for regular enrolled students. Hence, MOOCs offering authentication by SSO are not really open. Another drawback is that there are multiple SSO providers available. Cooperation among them seems unlikely and thus, the student might not have access to all MOOC providers. Furthermore, this approach might authenticate the student, however, it has not been designed to prevent cheating: Account sharing is still possible and it is not possible to check whether the authenticated student actually dealt with the tasks. Nevertheless, it is remarkable that this approach is usually free of charge.

Online Proctor Services are similar to the proctored exams approach presented above. In detail, these companies offer exam supervision through a webcam. Before taking the exam, the student is identified with his ID. Then the student is monitored while taking the exam [16]. Enhanced solutions broadcast the student's screen to the proctoring service as well in order to confine cheating. ProctorU[9] is one example. It offers proctoring services to various universities and the MOOC provider Udacity. The price is 15 $ per hour.

Online proctor services seem to be a real alternative to proctored exams, because this enables the exam to be taken from almost any location. However, it requires a lot of personnel, because every student has to be supervised by a single proctor. Assuming the used software is secure, this approach is the most suitable one to prevent cheating. However, due to the price, it is only feasible for exams and not for frequent use. Hence, the course assignments can still be handed in by a different participant than the student taking part in the exam. Another advantage of this approach is that the proctor can even offer individual help to the student.

Most current approaches still do not authenticate the student's identity effectively. Real-world identity proof is only guaranteed in proctored exams and not throughout the course. Hence, any other assignment is prone to either account sharing or at least collaboration. However, these are problems regular university courses have to deal with as well. Nevertheless, the physical presence of (non-online) proctored exams is a major drawback.

The other providers, presented in Section 2.1, have not relied on a special authentication approach so far. Nonetheless, Eliademy, Khan Academy and Udemy offer a digital certificate of participation; FutureLearn offers a hardcopy for a fee. P2PU pursues a different approach. The participants are encouraged to track their invested time and submit the overall time spend. Then, they obtain official certification of the numbers of hours invested. Furthermore, they proposed that all participants could sign each others certificates to convey credibility. Internally, the students can use a badge system to present their achievements to other students [3]. The provider MOOC.fi does not offer any certificates at this point. A summary about advanced authentication mechanisms and the possibility to obtain a certificate is given in Table 1.

[9] http://www.proctoru.com/

Provider	Adv. Authentication	Certificate
Coursera	Signature Track	digital
Udacity	Proctored Exams Online Protor Services	digital
Udemy	-	digital
edX	Proctored Exams	digital
FutureLearn	-	hardcopy
Khan Academy	-	digital
P2PU	(Single-Sign-On)	selfmade
Eliademy	-	digital
Mooc.fi	-	-

Table 1: Current Authentication Approach used by Presented MOOC Providers

Overall, all providers rely on login credentials which consist of email-address and password authentication. Most providers do not require the student's name during the registration process but instead a user-chosen display name. The student has to have access to the email-address, because the providers require the students to verify it. Other than that, no special security measures are taken.

3.2 Upcoming Approaches

This section introduces and evaluates newly proposed approaches for authentication in MOOC. The included approaches differ extremely: some are only theoretical (e.g. Multi level authentication), some have been used in other areas (e.g. Queries on Personal Identifiable Information), and some others exist only in a proof-of-concept similar implementation (e.g. Stylometry). Unfortunately, there is no solution which is suitable to solve the existing problems completely.

Biometric/Voice authentication has been suggested to identify students. This is a similar approach to Coursera's Signature Track. In the proposed scenarios, the authentication relies on (unique) biometric features to authenticate the student. Likely features for authentication include fingerprints, the student's face, voice recognition and the already deployed typing pattern [2]. However, for these authentication features the problem of initial registration respectively identification exists. For example, comparing the student's picture with his ID card's picture, could be setup for face recognition. Again, this proposed approach is only feasible, if every participant has access to a microphone, a fingerprint reader, or a camera respectively. Another challenge is the processing and storing of private data and information. Some students might not be content with sharing this amount of personal identifiable information (PII).

Until now, there have not been any serious attempts to offer authentication based on biometric features other than the typing pattern. With the increasing availability of the required sensors, this is expected to change. A likely scenario would include the student's smartphone into the authentication process, because it usually includes all necessary sensors.

Stylometry pursues the idea of identifying students based on their linguistic skills throughout the course. Hence, the student's submissions are constantly compared to a known linguistic style. Features for classification include word length, vocabulary richness, word shape, and frequency of characters and function words. In general, this approach suffers from the same problems as the previously introduced Signature Track: Initial identification and a text source for indexing a linguistic style are the main open challenges.

Currently, Narayanan et al. [13] claim that they are able to narrow down possible writers based on a single text of each writer. With an increasing amount of text, the probability of identifying individuals increases. Nevertheless, they are aware that the current approach is not feasible yet to provide authentication in MOOCs. For that reason, they recommend to link their proposed approach with another form of authentication. Additionally, they recommend follow-up evaluations to improve their classifier. Overall, the effect on privacy remains to be evaluated on a large scale. They warn that due to the gathered features, unintended application might offer the possibility to identify the author of nearly every published text [13].

Queries on PII as an approach has not been evaluated in context of MOOCs yet and it is unlikely that it will be applied in the future. Basically, the idea is that the student authenticating, has to answer questions only the registered student could answer [6]. However, this form is considered to be insecure, because nowadays, PII can be found in social networks such as Facebook or Twitter. Another problem for the providers is how to generate an initial "database" of questions to base the authentication on: If the student has to provide this information, there is no identity check performed. Furthermore, this approach is rather privacy invasive. If it would still be considered secure, the queries would have to include information the student would not share under usual circumstances. This raises the question, why should the MOOC provider get access to it.

Peer supervision is related to online proctoring services. Each student is assigned a peer and it is expected that the assigned partners authenticate each other. This approach is mainly based on trust and might work, because the peers are assigned randomly. Myers [12] even expects learning benefits from this approach due to the encouraged knowledge exchange. Unfortunately, it is unlikely that official institutions will accept this trust based approach as it is prone to cheating. Another drawback is that the peers have to study for the MOOC at the same time to authenticate each other. A solution is to introduce multiple peers. However, this increases the probability of cheating.

Multi level authentication aims to improve the authentication by offering a set of approaches for different situations. It is not a method independently supporting authentication. The provider should adapt the level of security based on the accessed resource. For example, it is more important that the student is effectively authenticated when taking an exam than when submitting an assignment. Hence, the authentication approach is assigned with respect to the use case. The reason for proposing this approach is that secure authentication might not offer sufficient usability to the student. Nonetheless, Miguel et al. [11], who proposed this concept, even suggest that multi-factor authentication might be introduced to combine multiple weaker approaches. They believe that alternating authentication methods helps to identify unintended access or even cheating.

However, there is no implementation available yet. Furthermore, the interaction between different authentication methods remains to be evaluated. Hence, it is difficult to estimate, whether this concept is feasible for advanced MOOC authentication.

Another proposed approach is described in Patent US20140157371 A1 [9]. The authentication should be performed based on the student's metadata. It describes how a provider can setup a secure testing environment bundled with student authentication. They attempt to identify the student with images as well. Additionally, they monitor browser metadata and other aspects to detect cheating. Unfortunately, a publicly available implementation is not available. Nevertheless, patenting shows how important solving the issue of insufficient authentication is.

4 Discussion

As seen in the previous sections, most approaches are insufficient for effective authentication for MOOCs. We identified the most important aspects which can be used to evaluate each approach: trade-off between security and usability, privacy concerns, inaccuracies, time of authentication, and costs. Currently, there is no solution available that satisfies each of the five aspects. Nevertheless, the most important constraint regarding MOOC authentication is the security [4], which is (at least indirectly) included in all presented aspects. Next, we go through each aspect in more detail.

Trade-off: Security and Usability The first aspect we identified deals with the trade-off between security and usability. The main questions that have to be answered are as follows. How much time and effort is the student willing to spend for authentication; How much information is the student willing to share with the provider or the authentication authority. If the security is too weak, authenticating the student is useless, because it is simple to circumvent the security. If the security is too strong, the student might be discouraged to participate in the MOOC. The invested effort might not be worth the participation. The providers need to find a balance to encourage the students to authenticate themselves, while not being able to compromise the approach. We will cover the provider's costs of storing and obtaining information separately in a following paragraph.

Privacy Concerns Another important aspect, which was already introduced in the trade-off between security and usability, is the amount of personal identifiable information that the approach requires. Basically, the question is how much user data has to be gathered to provide the authentication method and how much data the student is willing

to share. In that regard well-educated students might be more strict with the sharing that kind of information with the provider. Especially in the case of "free" MOOCs, the providers might try to increase their revenue by utilizing privacy sensitive information. Related is the question how long the provider stores the data and whether it will be used again if the student is participating in another MOOC of the same provider. Furthermore, the student might not be willing to share his PII to a third party, which is responsible for the authentication, courses with online proctoring currently require. As seen before, there are even approaches available, for example Stylometry, which might be able to identify the student in the Internet.

Inaccuracies Related to the security properties of the approach, are the inaccuracies it features and this is a vital criterion. Most solutions compare data of the current session to information that was gathered in earlier sessions or to documents that have been issued in the past. The crux is at which point the used approach is too inaccurate to lead to a rejection of MOOC certificates. If an used approach is too inaccurate, the usefulness of the issued certificates will decrease signaficantly. Furthermore, students might switch to another provider, if, occasionally, the authentication system does not recognize them. Therefore, the provider has to evaluate how error-prone the applied authentication mechanisms are for the sake of both legitimacy and usability.

Time of Authentication Time of authentication describes at which stage of the MOOC, the authentication and identification of the student takes place. There are approaches which are feasible for course-long authentication, while others might only be used for authentication before an exam. The provider has to decide when the authentication and identification should take place. Usual stages for that are for example upon course enrollment, before the exam, or for each course submission. However, this mostly depends on the used technology: Some approaches have a learning curve, hence, they need to be used from the outset on. Other approaches are too costly too be used frequently.

Costs The cost aspect is important for both the student and the MOOC provider. The main question is how much are they willing to pay for proper authentication and MOOC certificate acceptance. Non-profit providers might be able to offer the same authentication for a lower fee than commercial ones, because they do not have to add a profit margin. However, the students might still be discouraged to participate if the fees are too high. The drop-out rate of MOOCs shows that not every student that enrolled finishes it. For that reason, providers started to offer authentication as an additional service, which the student can pay if desired. Related to that aspect is the distinction between variable and fixed costs: Does a provider have to charge a student who already took an authenticated course the fixed costs again or does the student only have to pay the variable costs. In general, the providers have to keep in mind that introducing any kind of fees, decreases the number of likely participants. Furthermore, fees introduced by providers are not the only costs for the students, but also the costs for hardware that is required for the authentication, e.g. a webcam, a microphone, or a fingerprint reader.

Our identified aspects for effective authentication in MOOCs cover all areas that are important for the stakeholders. While privacy concerns mostly affect the students, the time of authentication is important for the authorities accepting the certificates. The costs and the trade-off between security and usability are important for both students and providers. Finally, all stakeholder groups are affected by the inaccuracies a solution might expose.

5 Conclusion

As we have seen in the presented approaches, there currently is no solution available that satisfactorily solves the authentication problem in MOOCs. Most current approaches convey security to the stakeholders. Even most of the proposed solutions feature drawbacks, thus they are not really feasible for guaranteeing the student's identity. Either the solutions are not usable, or too inaccurate, or too insecure, or just too costly. Another important aspects deals with the time of the student's authentication, i.e., which part of the course is (effective) authenticated. Furthermore, students might be discouraged to participate, because the deployed approach is too privacy invasive.

It is beyond dispute that authentication for students of MOOCs is an important factor. It remains to be seen how the acceptance rates of MOOC certificates will develop with the presented approaches in the future or whether one of the broad number of MOOC providers manages to create an entirely new concept. Authentication in MOOCs is a research topic that currently many scientists look into. As of today, the major research breakthrough still has not occurred. Surely, official acceptance by universities or companies will increase the boom of MOOCs even more: MOOCs can generate plenty economic value and revenue.

References

[1] N. Anderson. Moocs - here come the credentials. Retrieved on February 8, 2015. http://www.washingtonpost.com/blogs/college-inc/post/moocs--here-come-the-credentials/2013/01/09/a1db85a2-5a67-11e2-88d0-c4cf65c3ad15_blog.html, 2013.

[2] P. Bond. Biometric authentication in moocs. *Bachelor Thesis, University of Amsterdam's Digital Academic Repository*, 2013.

[3] A. Cole. Issue certificates. Retrieved on February 8, 2015. http://wiki.p2pu.org/w/page/32589272/issue%20certificates, 2011.

[4] P. Davidson and K. Hasledalen. Cyber threats to online education: A delphi study. In *Proceedings of the*

2nd International Conference on Management, Leadership and Governance: ICMLG 2014, page 68. Academic Conferences Limited, 2014.

[5] T. Haider. A comprehensive list of mooc (massive open online courses) providers. Retrieved on February 8, 2015. http://www.technoduet.com/a-comprehensive-list-of-mooc-massive-open-online-courses-providers/, 2013.

[6] P. Hyman. In the year of disruptive education. *Communications of the ACM*, 55(12):20–22, 2012.

[7] C. Inc. Introducing signature track. Retrieved on February 8, 2015. http://blog.coursera.org/post/40080531667/signaturetrack, 2013.

[8] M. Krause. A behavioral biometrics based authentication method for mooc's that is robust against imitation attempts. In *Proceedings of the first ACM conference on Learning@ scale conference*, pages 201–202. ACM, 2014.

[9] V. Le Chevalier and C. F. Geiger. Authenticated access to accredited testing services, July 3 2013. US Patent App. 13/935,150.

[10] A. Maas, C. Heather, C. T. Do, R. Brandman, D. Koller, and A. Ng. Offering verified credentials in massive open online courses: Moocs and technology to advance learning and learning research (ubiquity symposium). *Ubiquity*, 2014(May):2, 2014.

[11] J. Miguel, S. Caballé, and J. Prieto. Providing information security to mooc: Towards effective student authentication. In *Intelligent Networking and Collaborative Systems (INCoS), 2013 5th International Conference on*, pages 289–292. IEEE, 2013.

[12] D. Myers. Informal learning and authentication. Retrieved on February 8, 2015. http://www.mivu.org/About-Us/Leadership-Blog/ID/843/Informal-learning-and-authentication, 2014.

[13] A. Narayanan, H. Paskov, N. Z. Gong, J. Bethencourt, E. Stefanov, E. C. R. Shin, and D. Song. On the feasibility of internet-scale author identification. In *Security and Privacy (SP), 2012 IEEE Symposium on*, pages 300–314. IEEE, 2012.

[14] M. Parry. edx offers proctored exams for open online course. *Chronicle of Higher Education*, 6, 2012.

[15] L. Persen. Single sign-on. Retrieved on February 8, 2015. http://scandinavia.wiki.nmc.org/Single+Sign-On, 2014.

[16] C. Sandeen. Assessment's place in the new mooc world. *Research & Practice in Assessment*, 2013.

[17] C. Sandeen. Integrating moocs into traditional higher education: the emerging "mooc 3.0" era. *Change: The Magazine of Higher Learning*, 45(6):34–39, 2013.

How dense are cell towers? An experimental study of cell tower deployment

Ashok Rajendran
Aalto University School of Science
ashok.rajendran@aalto.fi

Abstract

In the past five years, there has been an amazing growth in the computational capability of smart phones when compared to the phones ten years ago. This rapid development has impacted mobiles with respect to power consumption and performance because the capacity of batteries has failed to keep pace with the other advances. To overcome these constraints, much research was performed to re-design the mobile applications such that power consumption is minimal. Cell-id positioning is one such advancement in location positioning, where location is determined based on a phone's signal strength vis-a-vis nearby cell towers. For efficient cell-id positioning implementation, we first need to analyze the signal strength distribution in a particular area to find the density of cell towers. We deployed a platform to collect cellular network information, including cell-id and signal strength, and we gathered data from different locations of the Helsinki area. This paper analyses the collected data and compares the density of cell towers in different areas of Helsinki. Also, it discusses several interesting phenomena observed from the data with respect to cell tower deployment.

KEYWORDS: cell-id; positioning; cellular network;

1 Introduction

In recent years, the functionality of mobile phones has increased rapidly. While still primarily intended for communication purposes, phones also possess PC-like capabilities. Various functionalities of a smart phone, such as voice communication, SMS and media playback, consume battery power. This makes the need for efficient energy management essential and it has prompted academia and industry to explore ways to improve the energy efficiency of mobiles. Research covers important aspects such as improving the transmission efficiency for wireless interfaces, designing smart interaction way and ingenious algorithms to reduce the working time of screen, CPU and sensors, and inventing new materials and power saving components [18].

A recent study by Carroll and G Heiser [1] reveals that wireless communication interfaces, mainly the interface of cellular network, contribute considerably to the overall power consumption of a mobile. According to Schulman et al. [13], power consumption of cellular network component depends on the signal strength. A poor signal consumes more power and vice versa. This has led to the development of many applications based on signal strength, which in turn consumes less power in mobiles. One major application that is considered is mobile positioning.

A positioning system is used to find the location of a mobile and having found it, many location based services are provided to it. A range of positioning systems exists, such as indoor and outdoor, depending on their usage scenarios. Indoor positioning system usually exploits Wi-Fi access points and sensor networks found inside the building to locate the mobile. Outdoor positioning locates mobiles in outside environments and based on the underlying techniques used, they are divided into three categories: satellite based, triangle based, and Cell-id aided positioning systems.

1.1 Satellite based positioning system

In satellite based positioning system, the mobiles are located in outdoor environments based on the signals acquired from the satellites. Mobiles need a specialized radio receiver to receive, recognize and analyze the satellite signals. This method can cover most of surface area of Earth, as well as near-Earth space [18]. Few of the current satellite positioning systems are Global Positioning System(GPS), GLONASS [10], Galileo [3], and Beidou navigation system. Among these, GPS is fully operational and widely used by many people. With a GPS receiver enabled in phone, a user will obtain his longitude and latitude after acquiring information from more than four satellites. GPS has advantages such as high accuracy where the accuracy of positioning varies from several meters to tens of meters. Thus, the GPS receiver is commonly used in aircrafts, navigators, mobile phones, tablets and wearable devices. However, it also has some weaknesses. First, users have to wait for 10 to 30 seconds before they acquire co-ordinates from satellites. If the environment is complex with many obstructions between the device and satellites, then the initial searching process will be even longer. Second, GPS receivers are battery intensive and consume more power when enabled. Many experiments revealed that the GPS component in smartphones consumes more than 300mW power [4] [9]. Thus, this method presents a problem for mobiles due to their limited battery capacity.

1.2 Triangle based positioning system

In triangle based positioning system, the mobiles are located by estimating the distance between the mobile and three nearby base stations. Using these values, triangle algorithm

locates exact position of the mobile. This positioning system can be classified further into three methods namely, time of arrival (ToA) method, angle of arrival (AoA) method, and time difference of arrival method (TDOA) [18]. Existing observed time difference (E-OTD) [14] is one of the TDOA methods which is widely used in US's enhanced 911 (E911) service Phase II implementation [2] for GSM network. In triangle based positioning system, the accuracy of distance estimation is affected by many factors such as multipath fading and channel conditions. This system has some disadvantages including low accuracy in positioning, the need for extra infrastructure for deployment and many more. These problems reduce the usability of this system.

1.3 Cell-id aided positioning system

Cell-id positioning locates position of the mobile with the help of cell id from the received signal. This method was recommended by 3rd Generation Partnership Project (3GPP) [19] as it is energy efficient and also does not require extra infrastructure and modification in software. Jonas Willaredt et al.[17] analysed the standards and protocols used in cell-id positioning system. This paper shows that there is no need of extra infrastructure for implementing this positioning system. Thus, the implementation of this method is very cheap and does not require any further upgrades. Energy efficiency of cell-id based system is studied by Jeongyeup Paek et al. [11]. A prototype of cell-id based system is developed and tested in this paper. The result reveals that this system consumes less power than GPS positioning system. However, the accuracy of the Cell-id based method is strictly based on the size of cell sectors. Stephan von Watzdorf et al.[16] studied the accuracy of different positioning systems using smartphones. Their study revealed that the accuracy of positioning in cell-id based system is above 500 metres whereas the accuracy is below 300 metres for GPS system. Thus, the cell-id positioning system works efficiently in small cells and not in large cells. This system is best applicable in urban areas as it has many small cells.

Due to its energy efficient approach in positioning a mobile and its applicability in urban areas, we analyse this system further. The received signal strength plays an important role in this system to locate a mobile in a particular cell. Received signal strength depends on the deployment of cell towers and the density of cell towers at a particular area. This paper analyses the density of cell towers and their deployment by collecting signal strengths from different locations of the Greater Helsinki area and performing a comparative study for a deep understanding.

The rest of the paper is organized as follows. Section 2 describes the signal collecting platform that was used in this experimental study for collecting signal strength. Section 3 describes the experiments that were performed in Helsinki city center specifically Kamppi, whereas section 4 discusses the results of the experiments and the interesting phenomena observed in the cellular network. Section 5 compares the results obtained from the Otaniemi and Kamppi city centre and the last section discusses future work in signal collecting platform and cell-id positioning.

Figure 1: App for collecting signal data

2 Signal collecting platform

The density of cell towers in a given location can be calculated by collecting signal strengths from mobiles at that location. To collect the signal strength from the mobile, we need a mobile application. An Android application has been developed by Jun Wu [18] for this purpose. Figure 1 illustrates the app which collects data by a process called wardriving [8]. War driving is the act of searching for the Wi-fi wireless network by a person in a moving vehicle using a portable computer or smart phone. Wardrivers use a Wifi-equipped device together with a GPS device to record the location of wireless networks. The results are then mapped with respect to its geolocation in Google maps. We used the same approach in our experiment where we collected cellular network information instead of Wi-fi networks.

A signal collecting app collects data by enabling GPS in mobiles to obtain a reference position. Then, it updates real-time cellular network information (cell id and signal strength) per one second. Furthermore, this application stops collecting data when GPS becomes invalid e.g, walking into indoor environments [18]. In addition to cell id and signal strength, this app also collects data such as latitude, longitude, and time stamp. The list of data collected are shown below.

- ***time*** : timestamp while collecting sample;
- ***latitude*** : latitude value in WGS84 standard;
- ***longitude*** : longitude value in World Geodetic System established in 1984 (WGS84) [5];
- ***gpsAltitude*** : height (h) above the reference ellipsoid that approximates the earth's surface [6];

```
time,latitude,longitude,gpsAltitude,accuracy,pressure,altitude,cellid,lac,ss,neighbors
1419954441168,60.16338257,24.92166482,90.0,36.0,1010.2644,24.884935,0,29000,-65,[]
1419954441187,60.16338257,24.92166482,90.0,36.0,1010.2644,17.808926,0,29000,-65,[]
1419954441198,60.16338257,24.92166482,90.0,36.0,1010.2644,17.808926,0,29000,-65,[]
1419954441999,60.16338257,24.92166482,90.0,36.0,1010.25635,17.877625,443042,29000,-65,[]
1419954442993,60.16346035,24.92163079,88.0,32.0,1010.2981,17.528845,443042,29000,-65,[]
1419954443993,60.16353705,24.9216006,86.0,36.0,1010.229,18.10486,443042,29000,-65,[]
1419954444994,60.16357622,24.92158903,85.0,36.0,1010.23535,18.052015,443042,29000,-65,[]
1419954445998,60.1636047,24.92159035,84.0,38.0,1010.292,17.579048,0,29000,-65,[]
1419954446999,60.16361905,24.92159347,83.0,45.0,1010.29517,17.552626,443042,29000,-65,[]
1419954447994,60.16362723,24.92159831,82.0,51.0,1010.30225,17.494495,443042,29000,-65,[]
1419954449011,60.16363532,24.92160614,81.0,52.0,1010.30225,17.494495,443042,29000,-65,[]
1419954449999,60.16363889,24.92161182,81.0,55.0,1010.2991,17.520918,0,29000,-65,[]
1419954451002,60.16364255,24.92161757,80.0,58.0,1010.3699,16.929049,0,29000,-65,[]
```

Figure 2: Sample records

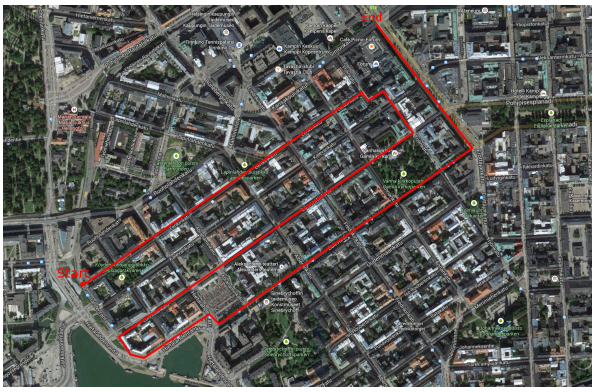

Figure 3: Path of our experiment

- ***cellid*** : unique number used to identify each Base transceiver station (BTS);
- ***lac*** : location area code;
- ***ss*** : average signal strength;

When the application is activated, it collects the above cellular information and updates it in a file. This file is in CSV format and it is stored in the file path allocated for this application. Each row in the file is called a record. Figure 2 shows the sample records in a file. At a particular location, if a device can connect to more than one base stations, then corresponding number of records are generated in the file. Whenever the application is switched off, it stops updating the file. Later, this file is copied from the mobile and processed to determine the density of cell towers.

3 Experimental design

After the implementation of the signal collecting application, we started to collect the data in Helsinki. We already had results analyzed from Otaniemi, which is suburban area of Espoo. In this experiment, we collected the data from Helsinki city center and compared it with our previous results. Our motivation was to find the distribution of cell towers in the urban areas of Helsinki where the population is dense. Therefore we chose areas near Kamppi, which is the city center. Figure 3 shows the streets and pedestrian routes covered by our experiment.

To get accurate results, we chose single mobile operator for our analysis. Operator Elisa was used throughout our experiment to collect data. We gathered datasets by performing

Phone	Model	Operator	Records
Samsung S4 with 3G	GT I9505	Elisa	3129
Samsung S4 with 3G	GT I9505	Elisa	3104
Samsung Galaxy Nexus with 3G	Galaxy_nexus	Elisa	8206
Samsung S4 with GSM network	GT I9505	Elisa	3066

Figure 4: Statistics of collected data

experiments using different mobiles many times. These mobile phones are Samsung Galaxy Nexus, Samsung Galaxy S4. Figure 4 shows the specifications of these mobiles. We also considered the network type used by phones while gathering our data. Records are collected separately using 3G network and GSM network.

4 Results and Observation

We collected totally 17,505 records at the end of experiments from different mobile phones. Figure 4 shows the exact number of records collected from each mobile phones. Analysis of these data showed some interesting phenomena with respect to the density of cell towers. The experiment was performed at Kamppi city center area whose diameter is approximately 2 kilometers. However, in this small region we detected nearly 74 different cell-ids for a single operator, Elisa. If the samples of all mobile operators are considered, then the total number of base stations deployed in the city center will be even more. Generally, it is considered that the cell size is counted in kilometers but the above result shows that there are 74 base stations within this small area. This reveals that the density of cell towers is high in urban areas.

4.1 Distribution of base stations

We analyse the results further to find each Cell-id's coverage, which reveals more interesting facts. For a given cell-id, the number of records found in the dataset represents the "perceived coverage area" of that cell-id [18]. Figure 5 shows the number of records and base stations on the X-axis and Y-axis respectively. The number of records illustrate the total number for the amount of base stations on the Y-axis. The graph depicts the size distribution of Cell-id's "perceived coverage area". Further analysis of graph shows that half of base stations at Kamppi cover fewer than 200 points in the dataset. However, there is also one large cell which has approximately 800 records.

4.2 Small cells and large cells

Cells can be divided into five types based on their dimension, namely nanocell, picocell, microcell, small macrocell and large macro cell. The size of these cells varies from some meters to kilometers. The size of a micro cell would be in the range of 0.1 to 1 Km whereas a small macro cell's range would be 1 to 3 Km [15]. Figure 6 shows one large and micro cell found in the Kamppi area. Measurements in the Kamppi area shows that most cells are either small macrocell or microcell. This observation reveals an important fact that

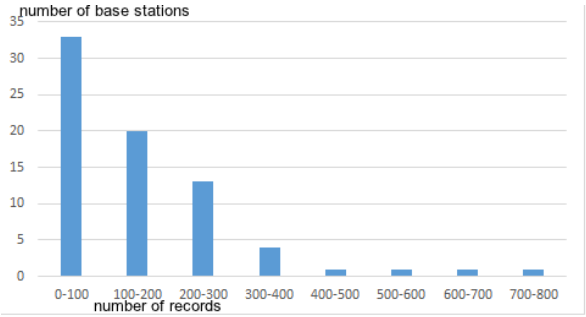

Figure 5: Distribution of base stations

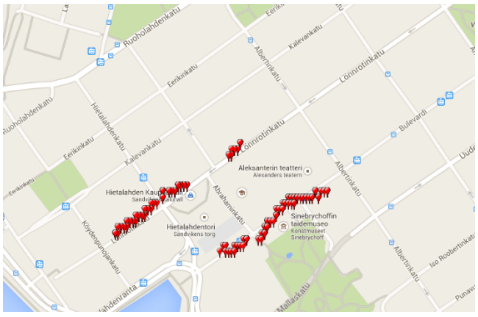

(a) large cell (396046)

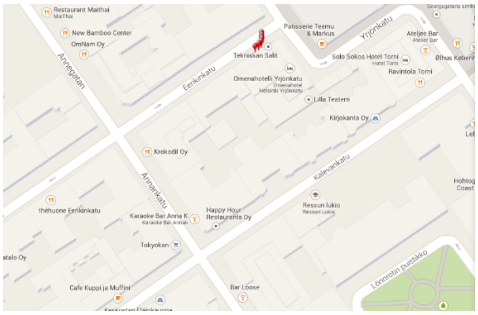

(b) micro cell (426737)

Figure 6: Types of cells

the area of coverage of each base station decreased a lot in urban areas. Reduced cell size improves the accuracy of positioning mobile phones in Cell-id based methods and also decreases positioning error [18]. In addition, we observed one more phenomenon where large cells are perceived as several small cells. There are many reason for this behavior. In this case, buildings block signals in the middle of the road, hence cell 396046 is divided and shown as three sections in figure 6a.

4.3 Intersection of multiple base stations

The third observation from the dataset reveals another fact: that a smart phone may connect to more than one base station at a given position. Figure 7 depicts the coverage of multiple base stations in the Kamppi area. The red point in this figure shows the position where only two cell-ids are detected. Green point represents the position where exactly three cell-ids are detected and blue point means more than

Figure 7: Coverage of multiple base stations

three cell-ids are detected at that position. It is observed that around 42 percent of the locations in the given area detected two base stations whereas remaining 58 percent locations detected three or more than three base stations. These results further confirm that a phone may connect to more than one base station. However, this observation contradicts the network planning theory. According to this theory, the cells are designed as hexagonal shape such that they use limited frequencies efficiently [12]. A base station is located in the center of each cell and so a smartphone will always connect to the closest base station to obtain the strongest signal strength. In other words, given a position, the optimally connected Cell-id is unique. However, our result showed that at a given position, a smart phone may detect more than one cell-id [18].

We also found that the base stations detected by the phones of a different network type differ. It means at particular location, the cell id detected by a 3G network phone is distinct from the cell id detected by GSM network phone. The collected datasets clearly show this result.

5 Comparison and Discussion

This section compares the datasets of Kamppi and the Otaniemi area. Similar experiments were performed in the Otaniemi area and the data were collected before as a part of the thesis work by Jun Wu [18]. These datasets are used for our comparative study. Figure 8 shows the path covered in Kamppi and Otaniemi. Both paths cover the same length of around 3.5 kilometers. For this same distance, the number of base stations detected in Otaniemi was around 67 [18] whereas in Kamppi, it was 74 cell-ids. This proves that density of cell towers at urban areas are slightly higher than suburban areas. However, the distribution of base stations are same in both Otaniemi and Kampii. Figure 9 shows the comparison of base station distribution in both Otaniemi and Kamppi. We can observe that more than half of cell-id detected at both areas, cover less than 200 points in the dataset. Additionally, most cells detected in these areas are either small macrocells or microcells. We also observe that a large cell is perceived as several small cells at both areas. Comparison of datasets also reveals the common fact that a smart phone may connect to more than two different towers at some locations.

This comparative study reveals the truth that cell ids sizes

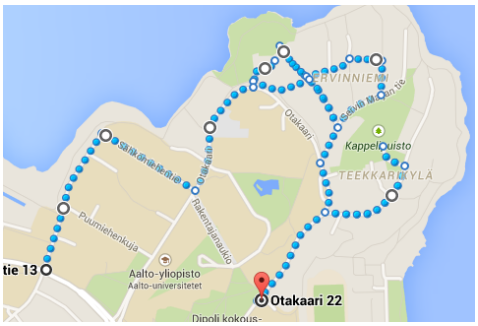

(a) otaniemi

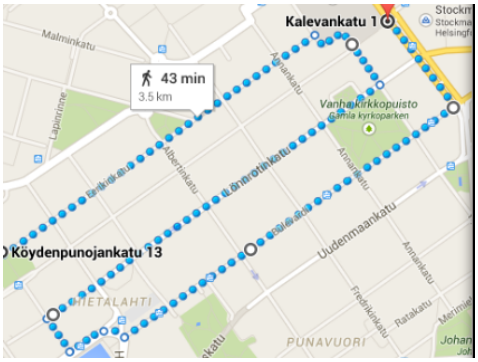

(b) kamppi

Figure 8: Paths covered in Kamppi and Otaniemi

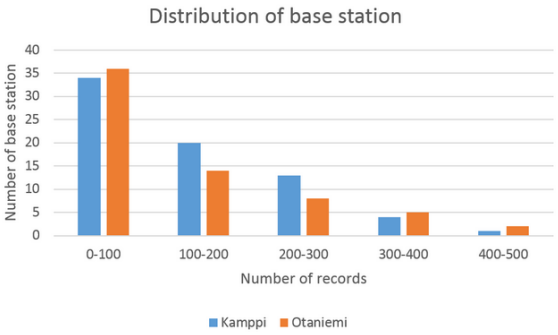

Figure 9: Distribution of base stations

are decreasing rapidly in urban areas. Cell size reduction helps in implementing cell-id positioning as the positioning error is less in small cells. Cell-id based method is energy efficient when compared to other techniques such as GPS positioning. Additionally, Jakob Hoydis et al. [7] proposed the concept of small-cell networks (SCNs) which is surrounded by the idea of cell-size reduction. Future telecommunication industry is also directed towards deploying dense self-organizing, low-cost, and low-power small cells to replace existing macrocells.

6 Future work

As of now, the experiments were performed in selected areas of Espoo and Helsinki with the operator, Elisa. Future work can be continued by performing following experiments, to obtain in-depth understanding of the density of cell towers. First, we need to collect data from all mobile operators and analyse it. Larger datasets provide more accurate results on how the cell phone towers are deployed at a particular location. A comparative study among the operators is also needed to analyse the characteristics of cell towers of each operator. Second, we need to perform experiments using different network types such as 3G, 4G, and GSM and compare the results. Third, it would be interesting if we could repeat the experiments using more smart phones and different models. This analysis would lead to more accurate results.

7 Conclusion

In this paper, we used the signal collecting platform developed by Jun Wu [18] and performed experiments in the Kamppi city center. The collected data were processed and compared with the results obtained in the Otaniemi area. The comparative study reveals many interesting facts about the deployment of base stations. The findings are summarized as follows:

- Density of cell towers are higher than expected in urban and suburban areas. We detected approximately 70 base stations in a small area whose diameter is 2 kilometers.
- Deployment of cell towers and their behavior were the same in both urban and suburban areas. We observed many macro cells and micro cells in these areas and concluded that cell sizes are being reduced in urban areas.
- Cell size reduction plays an important role in implementing cell-id positioning. This paves the way for energy efficient positioning systems in near future.

References

[1] A. Carroll and G. Heiser. An analysis of power consumption in a smartphone. In *USENIX annual technical conference*, pages 1–14, 2010.

[2] F. C. Commission. Wireless 911 services, June 2014.

[3] CONSORTIUM. *The galilei project: Galileo design consolidation*. European Comission, August 2003.

[4] I. Constandache, S. Gaonkar, M. Sayler, R. Choudhury, and L. Cox. Enloc: Energy-efficient localization for mobile phones. In *INFOCOM 2009, IEEE*, pages 2716–2720, April 2009.

[5] B. L. Decker. *World geodetic system. Tech.rep.* DTIC Document, 1986.

[6] W. Fraczek. Mean sea level, gps, and the geoid, July 2003.

[7] J. Hoydis, M. Kobayashi, and M. Debbah. Green small-cell networks. *Vehicular Technology Magazine, IEEE*, 6(1):37–43, March 2011.

[8] C. Hurley. *WarDriving: Drive, Detect, Defend: A Guide to Wireless Security*. Elsevier Science, 2004.

[9] M. Kjaergaard. Minimizing the power consumption of location-based services on mobile phones. *IEEE Pervasive Computing*, 8(4), 2010.

[10] W. Lechner and S. Baumann. *Global navigation satellite systems*. Computers and Electronics in Agriculture, 2000.

[11] J. Paek, K.-H. Kim, J. P. Singh, and R. Govindan. Energy-efficient positioning for smartphones using cell-id sequence matching. In *Proceedings of the 9th international conference on Mobile systems, applications, and services*, pages 293–306. ACM, 2011.

[12] M. Rahnema. Overview of the gsm system and protocol architecture. *Comm. Mag.*, 31(4):92–100, Apr. 1993.

[13] A. Schulman, V. Navda, R. Ramjee, N. Spring, P. Deshpande, C. Grunewald, K. Jain, and V. N. Padmanabhan. Bartendr: A practical approach to energy-aware cellular data scheduling. In *Proceedings of the Sixteenth Annual International Conference on Mobile Computing and Networking*, MobiCom '10, pages 85–96, New York, NY, USA, 2010. ACM.

[14] S. Tekinay. Wireless geolocation systems and services. *Communications Magazine, IEEE*, 36(4):28–28, April 1998.

[15] E. Trevisani and A. Vitaletti. Cell-id location technique, limits and benefits: an experimental study. In *Mobile Computing Systems and Applications, 2004. WMCSA 2004. Sixth IEEE Workshop on*, pages 51–60, Dec 2004.

[16] S. von Watzdorf and F. Michahelles. Accuracy of positioning data on smartphones. In *Proceedings of the 3rd International Workshop on Location and the Web*, page 2. ACM, 2010.

[17] J. Willaredt. Wifi and cell-id based positioning-protocols, standards and solutions. *SNET Project WT*, 2011.

[18] J. Wu. Signal Collecting Platform and Hand- print Positioning System. Master's thesis, Aalto University, June 2014.

[19] Y. Zhao. Standardization of mobile phone positioning for 3g systems. *Communications Magazine, IEEE*, 40(7):108–116, Jul 2002.

User Authentication or Identification Through Heartbeat Sensing

Sowmya Ravidas
Student number: 451125
sowmya.ravidas@aalto.fi

Abstract

One of the most common methods for authenticating people are passwords. However, passwords are difficult to remember and are also susceptible to attacks. An alternative for password is the traditional biometric system, for example, fingerprints. Although such biometrics provide better security compared to passwords, these are not suitable for people with disabilities. In this paper, we discuss ECG based authentication systems that use heartbeat for authenticating people. Compared to other biometrics, ECG based authentication systems are secure as heartbeats are difficult to forge and are also suitable for people with disabilities. Hence, such systems can be a potential alternative for traditional biometrics. In this paper we discuss the process of measuring ECG, explain the workings of ECG based authentication systems and also discuss a commercial product that uses this technology. The main contribution of the paper is our analysis on the suitability of such systems for the public and we also discuss the challenges associated with the same.

KEYWORDS: Heartbeat, Hear-rate Authentication, Identification, ECG.

1 Introduction

Passwords are a common method for authentication and identification of people. We use atleast 4-5 passwords a day for various purposes including logging into our computers, email, opening electronic door locks etc. One problem for users is remembering these passwords; a second is their vulnerability to attacks such as intrusion attacks, social engineering attacks and denial of service attacks.

One solution to both these drawbacks is Biometric authentication. This approach to authentication is achieved based on physical traits such as finger prints, eye scans and facial features. Though biometric authentication is expensive, it is widely being used in places such as passport offices and visa service centres. However, these authentication methods do not scale well for all users. For example, finger print based authentication is not a feasible method for physically challenged people and also for people with temporary cuts or wounds on their fingers. It is possible that finger prints fade away over time due to moisture level or due to diseases such as cancer. The current biometric systems requires the users to authenticate every time they use a service, which is time consuming.

This report focuses on using one's heartbeat information for secure authentication and identification. It has been found that the biometrics from the human heart such as heartbeat information show uniqueness, because it is inherited from the individuality of DNA (Deoxyribonucleic acid).

ECG based authentication methods could be one of the potential methods for efficient authentication. Fahim et al [19] describe how heartbeat measurement by ECG(Electrocardiogram) can be used as a standalone biometric system for authentication.

ECG measures the electrical signals of the heart via sensors that are connected to the body. From an ECG, we can measure the inter pulse interval, which is the time lapse between any of two nerve impulses and also measure the heart rate variability (HRV). HRV shows the beat-to-beat alterations of the heart rate. The inter pulse interval or the heart rate variability can be efficiently used to identify individuals. Another advantage is it is possible to collect the inter pulse interval from any parts of the body, which makes the process easier.

There are also many other ways to measure heartbeat information besides ECG. Balakrishnana et al [13], explain how to detect pulse measurements from the head motions of a person which they capture in a video. Further analysis of those revealed that the extracted pulse information shows similarity to the ones obtained using ECG.

In addition to the above, there are also several products that come with a heart rate monitor embedded in them which helps in measuring the pulse and which also give health advice to the users. Polar heart rate monitors [8] are an example.

Wearable devices for seamless authentication such as Nymi [6] is one of the latest developments in this field. Nymi is a wristband that uses heartbeat information to perform seamless authentication of people to devices in various environments. According to the researchers of Nymi, heart rate is unique even when we exercise [1]. With a click on the band, the users can authenticate to their computer, doors, cars etc and they are no longer required to memorize their passwords.

The remaining part of the report is structured as follows. Section 2 explains how ECG is measured. The subsequent section concentrates on authentication based on ECG and also discusses about workings of the Nymi band. We then look into some of the security challenges in using ECG based methods for authentication or identification. The last section concludes the report.

2 Measuring Heart-rate using ECG

In this section we look into measuring ECG, which is traditional but a widely used method in hospitals to diagnose a cardiac patient. An ECG records the electrical activity of the heart by using sensors placed on the body to detect these electrical changes. Usually 10 sensors are placed on the limbs and 2 on the chest. The heart's potential is measured in different angles and are recorded as a time series graph.

The electrocardiogram report consists of this graph plot of voltage versus time, which shows the overall magnitude of the heart's electrical potential. This gives accurate information about heart rate and rhythm.

Our human heart consists of chambers- Left Atrium, Left Ventricle, Right Atrium and Right Ventricle. The atrium and ventricle contract and relax together in-order to pump blood throughout the body. There is an electrical system that makes this possible by creating electrical impulses. The impulse occurs in cells located in the right atrium and spreads across the walls of the atrium causing the contraction. These group of cells commonly known as SA node is responsible for setting the heart rate and rhythm. Each heartbeat is represented on the ECG as a PQRST complex, where each of these letters represents a part of this complex, as seen in Figure 2.

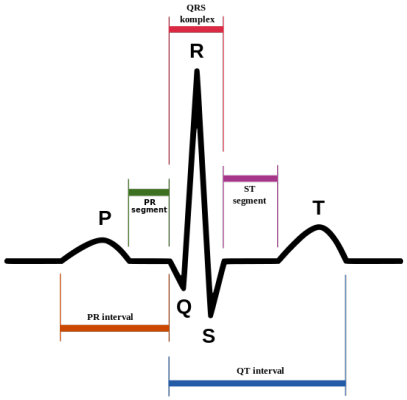

Figure 1: PQRST Wave [9]

1. P wave
   The P wave is the first part of the complex and it represents the atrial depolarization. This is followed by atrial contraction.

2. QRS complex
   The next part of the complex is the QRS wave, which represents the ventricular contraction. A negative deflection in the QRS complex is referred to as Q wave. A Q wave may or may not be present in an ECG. The R wave is the largest wave because that reflects the depolarization of the main parts of the ventricles. The negative S wave follows the positive R wave. This wave represents the final depolarization of the ventricles.

3. T wave
   The last component is the T wave, which represents the re-polarization of the ventricles. Usually, it is the T wave that provides most information about cardiac abnormalities.

If the P wave precedes the QRS complex, then the rhythm is said to be correct. If not, the person may be suffering from cardiac arrhythmia, irregular heartbeat or other cardiac problems.

There are also techniques that measure ECG from the finger tips and avoids placing them on the chest. Louren et al [16] found that only 3 leads on the index finger are enough to identify a person. This consumes less time and is also easier for the clinical experts to measure them.

However, these ECG measurements still requires contact with the body. Guha et al [13] showed that it is possible to measure heart rate and beat lengths by measuring head motions that are captured in a video. They extract interesting cardiac information from the head oscillation by applying filtering and PCA techniques.

In the next section, we discuss how ECG data is used for authenticating people.

3 Authentication and Identification through Heartbeat

This section discuss the ECG based Biometric system and how it enables secure authentication. We will also look in detail at a commercial product- Nymi and evaluate its security.

3.1 Uniqueness of Heart-rate

Heart-rate is unique and it differs from person to person. This uniqueness occurs because our heart inherits it via the DNA.

DNA gives distinctive properties that help in distinguishing people and also accurate identification of them. This genetic information flows from DNA to RNA to Protein and protein is responsible for the uniqueness provided by finger prints, Iris and other biometric data [19].

The shape of the heart, face and other organs also exhibits unique features that are derived from the individuality of the DNA and can be used for identification of people. ECG shows the electrical activities of heart and it can also be used for successful identification because it is, in turn, derived from DNA.

3.2 ECG based Biometric

The inter-pulse interval and the heart rate variability can be used to identify individuals.

This method was first introduced in Biel et al's paper [14] where they collected ECG of people and performed a multivariate analysis. Their results revealed the possibilities of using ECG as a biometric system.

Shen et al, [17] used a one lead ECG where the leads are kept on the hands and they further analysed the QRS complex and the T wave. Using template matching and neural

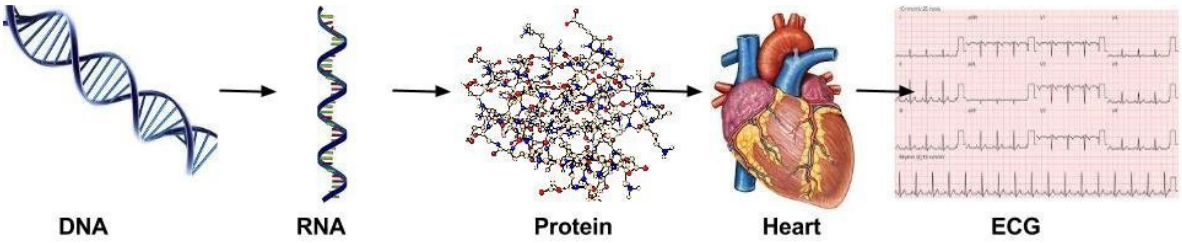

Figure 2: ECG Biometric from DNA

network classifiers they claimed to have received an accuracy upto 100% for 20 subjects.

These results shows the accuracy of ECG based authentication and also its ease of use. These methods are suitable for everyone including physically challenged people. Also, the heart rate can be measured from any parts of the body which makes the process much easier.

In an ECG based biometric system, an enrolment stage is conducted to collect a user's ECG and store it in a database. During the authentication stage, a one-to-one mapping of the ECG templates is performed. A person can authenticate himself with a PIN or a smart card which has his details such as name. The system then captures his ECG and is mapped against the ECG originally captured during the enrolment stage for this person. If there is a match, the authentication is considered to be successful.

On the other hand, identification can be achieved by performing a one-to-many mapping. The acquired ECG is mapped with all the ECG templates in the database. The individual can be identified if any of the templates are mapped with high probability.

During the enrolment stage, the unique feature of the ECG is identified and during identification, a pattern matching of this feature is compared with all the other ECGs in store.

3.2.1 Classification of ECG based Biometric

There are two kinds of features which can be extracted.

1. Time Domain Feature Extraction

   The unique features are extracted from the ECG and the feature wave duration, amplitude, direction and other details are obtained. These component values are saved during the enrollment stage. During authentication, all these data are extracted and a one-one template matching is performed.

   The features extracted can be further classified into two types:

   (a) Extracting Morphological Features

   This reveals the time domain features from the ECG, such as P wave duration and amplitude, QRS duration and amplitude and T wave duration and amplitude. The collected data is used for identifications although most of them are used traditionally for cardiac diagnosis.

   (b) Uniqueness of Beating Patterns

   Consecutive heartbeats can also reveal unique information such as heart rate variability or interval. These features depend on the breathing pattern, heart rhythm but can be used for identification purposes.

2. Frequency Domain Feature Extraction

   The ECG signal is converted from time domain to frequency domain. To extract the feature wave, transformations, such as Fourier transform, wavelet transform, discrete transform are applied.

In general it is difficult to achieve successful mapping due to the time-varying nature of ECG, as mentioned earlier. Abnormality in ECG or PQRST signature can result in erroneous template matching which can cause a denial of service. According to a few scientific works, a common situation where change in ECG happens is when the stress level of the users vary.

However, Israel et al [15], explained from the experiments, that the ECG signals does not really change under anxiety or stress. The authors simulated the 7 states of mental stress on people and were able to successfully verify their identity under varied stress levels. From the results, we can say that in most cases it is possible to identify an individual under varying stress less, but we should also keep in mind that it is not 100% accurate. Also, this experiment did not consider any physical activities of a person.

It is difficult to collect an ECG in a shorter time and also to provide unique features in real time. Hence, ECG based biometric system can take much longer time to collect than the finger prints or Iris data.

Also, there is not much experimental data available compared to the amount of finger-prints or Iris data. Researchers would need more ECG based data to analyse the system further and make it globally acceptable.

3.3 Nymi Band

Nymi Band is a product from the company- Nymi, based in Toronto. The idea behind this product is to perform seamless authentication in various environment using a single wristband. The band uses heartbeat information for authentication and hence the researchers at Nymi claims that it is not only easy to use but also highly secure. With a click on the band,

the users can authenticate to their computer, doors, cars etc and they are no longer required to memorize passwords.

3.3.1 Components

HeartID is the core of Nymi band and it helps users authenticate to the computers and other environments. HeartID is attached using sensors that record the heart's signature and confirms the identity of the individual.

There is a Nymi Band Core component that has the encrypted hardware that protects the communication. It also contains motion sensing components such as a gyroscope for gesture recognition and also a haptic feedback motor.

The LEDs on the Nymi Band increase it's usability. It notifies users about it's current state, battery life and other messages. The communication is performed using LED patterns.

Also, we can place the band on our wrist which increases the usability. If the band is cut, then the circuit is disabled which prevents from unauthorized access. The band detects that it has detached from the body and disables the system.

They also claim that since the band is required to be tied to the wrist, it is made of not so allergic materials which enhances the comfort of the users.

3.3.2 Security Analysis of Nymi

The authentication used in Nymi is 3-factor authentication, which consists of heartbeat, Nymi band and also a smart phone.

Figure 3: Three Factor Authentication [3] [4] [5]

1. Heartbeat based authentication
   The authentication is performed based on the collected ECG data. The ECGs are difficult to copy or forge. Hence, it makes it harder for the attacker because inorder to perform the attack, he has to simulate user's ECG.

2. The Nymi band
   If the band is broken and stolen, it cuts the circuit and the band stops working. Even if the band is stolen intact, the band detects its detachment from the body and it invalidates the authentication. Hence, attacker cannot gain access even if the band is stolen.

3. Smartphone
   Let us assume that the attacker has succeeded in spoofing the band. He would still be unsuccessful in gaining access. This is because, in addition to the band, a smartphone is required for successful authentication. The phone will have the Nymi application that enables pairing of the band and the authentication system. This is required for initial authentication and is not required on later stages.

It is difficult to capture user's heartbeat or ECG and hence it is not that easy for an attacker to replicate it. If the attacker breaks the band or steals it, the device stops working as there is no match between the heartbeat patterns. However, it is possible that the attacker can steal the smart phone while it is being authenticated. Since this is a 3 factor authentication, the attacker may also have to spoof the wristband in addition to the phone to be able to access the system.

According to Nymi [7], it is based on Heartbeat based authentication and hence it can be more secure than the traditional biometrics because it is difficult to forge someone's heartbeat.

Experts at Nymi says that Nymi prevents a user from impersonation attack and passive eavesdropping. It is not possible for someone to impersonate as the heartbeat is unique and are different for different people. Also, it is hard to eavesdrop as the band is always tied to the user's wrist and is difficult to eavesdrop without the knowledge of the user. They also claim that such features also prevents some of the active attacks such as man-in-the-middle attacks.

Nymi uses robust cryptographic techniques which guarantees that only the devices that are paired with Nymi can detect the presence of it. This prevents tracking by adversaries and also preserves privacy to certain extend.

4 Security Challenges in Using ECG for Authentication or Identification

4.1 Accuracy

It is said in [1] that the heartbeat based authentication is slightly less accurate than using fingerprints. This is because heartbeats vary continuously and is not consistent over time.

As mentioned in Section 3, a heartbeat based authentication is successful in most of the cases even under extreme stress levels. However these results are not very accurate and usually people take measurement when subjects are in idle state.

The heart-rate can vary to extreme cases and it might make it difficult for the system to identify the individual. For example, if a person who has survived a heart-attack, the heart rate would vary considerably from what it was when he actually recorded. The same might be the case for a person with breathing problem.

Factors including physiology, geometry of the heart, body build, gender and age makes ECG based authentication less accurate when compared to other biometric products. There are also other factors that change on a slow rate such as age, body habitus etc.

Hence, accuracy is one of the major challenges in such systems. There are some scientific works that talks about probabilistic authentication model based on the activities of the user [18]. Although the claim scales for entire world, the probabilistic model may not be a really scalable solution and in most of the cases the sample size they have considered is quite small.

4.2 Attacks

The ECGs are captured in the ECG biometric system and is transferred to ECG repository where it is stored. The ECG is transferred through public internet and it is very important that ECGs are encrypted before transmitting. If not, there are possibilities for spoofing attacks, where the attacker spoofs the heart rhythm and replay it to the machine, which can allow an attacker to access the captured ECG.

A sufficiently strong method of ECG encryption is by using Permutation Ciphers. Permutation ciphers are now being used by researchers to encrypt ECGs before sending it over the Internet. Here, the original ECG is encrypted by applying mathematical transformations, which results in random ASCII letters. The permutation key is known only to the receiver, who can decrypt the encrypted ECG. This technique is secure when compared to AES and DES according to [19]. However, a permutation cipher is prone to attacks and is easy to break if the attacker can reconstruct the initial shift key. If this method can be combined with existing encryption schemes such as AES and DES, the system will be more secure.

Other possible attacks could be denial of service, i.e if the attacker manages to break either of the authentication factors, they can cause a DoS attack. Let us take Nymi band as an example. It is possible to cause denial of service if the attacker manages to hack either the Nymi band or smart phone or the authenticating system. If the band is broken or stolen, the user cannot authenticate himself. It is also possible that the smartphone is stolen and the user is no longer able to access the server.

Another possibility is where the attacker makes an attempt to confuse the system resulting in gaining access for him or causing a DoS for the user. For example, if the attacker has user's personal data, he can feed his own ECG and try to gain access. This will of-course result in a failure because the ECGs do not match. However, if the attacker tries multiple times, the system may lock the user and he can no longer access it.

Social Engineering attacks are one of the highly possible attacks in ECG based authentication systems. Attacker can possibly trick the user to grant authentication for him. For example, the attacker can trick the user to gain his permutation key. This enables attacker to gain access to the original ECG, learn user's data and also his health information.

Attacks are possible from family members. For identical twins, the DNAs are similar and hence they have more chances of having similar ECG data.

Most of the current works related to ECG based authentication that we have seen lacks a proper adversarial model. It is necessary to have an adversarial model to study the system in detail and to identify and correct the flaws in the architecture.

4.2.1 Thread Model

In this subsection we design a thread model and analyse the security issues, attacks and motivation of the attackers.

Some of the important **assets** in ECG based authentication systems are:

1. Personally identifiable information (PII)
2. Heartbeat data
3. User credentials
4. Phone or other device
5. Measurement System
6. Health information
7. ECG mapping system
8. Money, if this authentication is used in banking systems.

Let us now evaluate the **potential attackers and attacks**.

1. Friends and Family: People trying to use other's ECG authentication. This is possible when users share their credentials with others and it is possible that they can take advantage of it. Family members for example, identical twins can have similar ECG patterns.

2. Criminals and Hackers: To steal information and make money. Criminals and Hackers can steal ECG information and try to feed fake ECG signals into the system. This causes an intrusion attack. Attackers can also hack the ECG mapping system and steal all the data which can reveal personal information and health information of the users. They can also steal a band or any other devices and cause a denial of service attack. It is also possible to copy the information and replaying it in another machine to gain access, try to do illegal things, to access prohibited computers or areas.

3. Insider attackers: These are attackers who are within the system. For example, Hospital authorities or people taking ECG measurements: They can record fake heartrate readings instead of original ones.

4. Companies who build products that measure heartbeat: These people can collect huge data without user's permission and use it for analysis.

   Even if the attack on gaining access from the ECG system fails, it is possible for the attacker to learn about personally identifiable information and track the user. Attackers will also get to know about the health status of the user. It is also possible that commercial products and companies can deny service at some point and demand more money from the user. User's health record can be used for other purposes without user's permission, which hinders user privacy.

   Having a good threat model in place helps to get a clear vision on the problems and focus on the parts that need improvement. The model helps to understand the motivation of attackers, prioritise the risks and solve them. It also helps in identifying more vulnerabilities, patch them and make the system more secure.

4.3 Privacy

In ECG based authentication methods, we don't really use private information such as photos or videos of users. Hence we can say that the privacy is preserved to certain extent. However, there are still concerns regarding privacy issues in using heart rate information for identity.

The main concern regarding privacy is with respect to who owns the data- the user or the server provider. The user's ECG and their personal details should be stored in a highly secure environment and must be sent through encrypted channels only. The service or product has to ensure that users can control their details and no one else is authorized to view or modify it. A proper agreement has to be made between the users and the service providers before the latter can use the information for other purposes.

ECG based data can reveal a lot on the person's health status. This can be private details and users may not want to share it with others.

Another concern with respect to privacy is how the ECG is measured. Users may prefer placing leads on the fingers than on the chest. Contact-less authentication explained in [13] is a good method, but users might not prefer it because it takes video of the user.

4.4 Usability

Usability is one of the major issues in ECG based authentication systems. If we use traditional ECG system to measure the ECG, users may find it difficult. This requires more time as users have to lie down with 12 leads placed on their body and sometimes doctors use gel to place these leads which can be inconvenient for the users. However, latest research talks about measuring ECG from fingers [16] and also measuring ECGs without contacting the user- such as from head motions as described in [13].

The commercial products such as Nymi comes with high usability and there is not much action required from the user side. However, always wearing a band can be inconvenient for the users. It is possible that they can forget to wear the band sometimes and for this reason they might want to shift back to passwords.

There are many scientific works that talks about accuracy and that ECG data is trustworthy. This makes it more usable because users will find it trustworthy as well as the system does not annoy them in providing the services. However, a matter of concern is how ECG changes over time, over several years. This is a problem that needs to be investigated further because if the ECG changes over time, it can have a direct impact on the usability.

Another concern is that users should experience that this method is secure. Sometimes it can happen that users does not really get the idea of what is happening and may feel that the system is not as much secure when compared to traditional password based system. In other words, users don't find a system to be secure until they type in a password.

There are advantages and also challenges with ECG based system. Hence, it is difficult to predict if this will be a widely accepted technology for authentication.

5 Conclusion

Passwords and traditional biometrics come with lot of disadvantages. The former is hard to remember and the latter is not suitable for all people. Also, both passwords and traditional biometrics such as finger prints are vulnerable to attacks. In this paper, we studied Heartbeat based authentication methods which provides a highly secure platform for the users to services. We looked to the various aspects of ECG based authentication and have also studied a commercial product. We discovered that ECG based authentication systems are suitable for all people and are secure. However, there are many challenges associated with the same. We looked into the accuracy, attacks, privacy and usability issues. We also designed a threat model that can help in evaluating and improving the security of such systems. ECG based authentication method is still an emerging field and we hope that it can provide the means to securely authenticate people.

References

[1] First look: Startup readies heartbeat-based authentication. WWW News article on Nymi: http://tinyurl.com/pofyq8b.

[2] Function of heart. WWW:Human Heart Anatomy http://www.livescience.com/34655-human-heart.html.

[3] Human heart. WWW ecg from heart: http://www.core77.com/posts/26369/forget-passwords-soon-your-heartbeat-may-open-your-e-mail-26369.

[4] Human heart. WWW ecg from heart: http://pixabay.com/en/cardiac-pulse-systole-heartbeat-156059/.

[5] Human heart. WWW ecg from heart: http://www.pcpro.co.uk/smartphones/6932/the-18-best-smartphones-of-2015-whats-the-best-phone.

[6] Nymi. WWW page of Nymi: https://www.nymi.com/.

[7] Nymi band white paper. WWW:Nymi Band White Paper https://www.nymi.com/wp-content/uploads/2013/11/NymiWhitePaper-1.pdf.

[8] Polar- hear rate monitor. WWW page of Polar: http://www.polar.com/en/products/.

[9] Pqrst. WWW:PQRST wave http://upload.wikimedia.org/wikipedia/commons/thumb/0/09/ECG-PQRST%2Bpopis.svg/800px-ECG-PQRST%2Bpopis.svg.png.

[10] Pqrst. WWW:PQRST wave http://www.emergsource.com/?page_id=90.

[11] Pqrst. WWW:PQRST wave http://www.ncbi.nlm.nih.gov/books/NBK2214/.

[12] Sa node. WWW:SA node https://my.clevelandclinic.org/services/heart/heart-blood-vessels/how-does-heart-beat.

[13] G. Balakrishnan, F. Durand, and J. Guttag. Detecting pulse from head motions in video. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 3430–3437. IEEE, 2013.

[14] L. Biel, O. Pettersson, L. Philipson, and P. Wide. Ecg analysis: a new approach in human identification. *Instrumentation and Measurement, IEEE Transactions on*, 50(3):808–812, 2001.

[15] S. A. Israel, J. M. Irvine, A. Cheng, M. D. Wiederhold, and B. K. Wiederhold. Ecg to identify individuals. *Pattern recognition*, 38(1):133–142, 2005.

[16] A. Lourenço, H. Silva, and A. Fred. Unveiling the biometric potential of finger-based ecg signals. *Intell. Neuroscience*, 2011:5:1–5:8, Jan. 2011.

[17] T.-W. Shen, W. Tompkins, and Y. Hu. One-lead ecg for identity verification. In *Engineering in Medicine and Biology, 2002. 24th Annual Conference and the Annual Fall Meeting of the Biomedical Engineering Society EMBS/BMES Conference, 2002. Proceedings of the Second Joint*, volume 1, pages 62–63. IEEE, 2002.

[18] J. C. Sriram, M. Shin, T. Choudhury, and D. Kotz. Activity-aware ecg-based patient authentication for remote health monitoring. In *Proceedings of the 2009 international conference on Multimodal interfaces*, pages 297–304. ACM, 2009.

[19] F. Sufi, I. Khalil, and J. Hu. Ecg-based authentication. In *Handbook of Information and Communication Security*, pages 309–331. Springer, 2010.

A Survey on Performance of Scalable Video Coding Compared to Non-Scalable Video Coding

Martijn Roo
Student number: 466686
martijn.roo@aalto.fi

Abstract

This paper presents a comparison of scalable video coding and non-scalable video coding based on current literature. The comparison focuses on properties of both coding techniques relevant to different use case scenarios. It concludes that SVC provides a better live-streaming video quality, a better storage efficiency, and better caching performance compared to AVC, whereas AVC decoders use less computational resources and video streaming services using AVC encoded videos have a lower operating cost when no caching servers are used. The operating cost of a video-on-demand service using caching with or without SVC is an area for further research.

KEYWORDS: video, encoding, mobile, multimedia, streaming

1 Introduction

Users connect to the Internet through various types of connections each having varying levels of available bandwidth. As a result, video streaming services have to take into account the heterogeneous connections with which their users connect to their services. They often accomplish this by using HTTP Adaptive Streaming (HAS). This works by splitting a video into small time-based segments, whereby each segment is made available in multiple qualities and therefore in multiple bit rates. This enables a client streaming a video to dynamically, during playback, choose in which quality it wants to download future segments based on the available bandwidth.

Figure 1 gives an overview of a HAS application. There is a regular HTTP server serving requests from a client device. The client device first requests the media's content description for the media it wants to stream. The client device continuously decides which quality segment to request next based on the available qualities, the available bandwidth, and the current buffer size.

Often HAS is used with different files encoded at a different bit rate to support different video qualities. This is non-scalable video coding and the server in Figure 1 shows this technique. Another option for supporting adaptive streaming is to encode a video file using scalable video coding in which case the file is encoded into layers and the number of requested layers determines the bit rate and quality of the video.

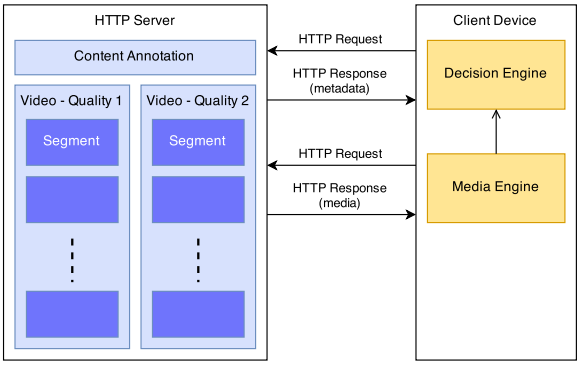

Figure 1: HTTP Adaptive Streaming architecture

Scalable and non-scalable video coding have different properties with regards to coding complexity, storage efficiency, operating cost, and caching performance. This paper compares HTTP Adaptive Streaming techniques using scalable and non-scalable video coding on these properties.

The remainder of this paper is structured as follows: Section 2 describes the background of HAS and SVC; Section 3 compares scalable video coding and non-scalable video coding on coding complexity, quality of experience, storage efficiency, and caching performance; Section 4 analyses the cost of hosting a video streaming service with SVC compared with AVC; finally, Section 5 and Section 6 conclude this paper and give recommendations for future work respectively.

2 Background

Applications of HTTP Adaptive Streaming allow the client to request a different video bit rate per fixed-time segment (usually around 2 seconds in length). Most HAS applications use non-scalable coding, which means that every video segment has to be encoded separately for every bit rate [5]. Usual methods for obtaining different bit rates are temporal, spatial, or quality scalability [14]. Temporal scalability allows creating a lower bit rate video by dropping complete pictures from the original bitstream. The resulting stream has a lower frame rate than the original. For spatial scalability, a video is coded at multiple picture sizes. Quality scalability is also known as SNR (Signal-Noise Ratio) or fidelity scalability and for this, a video is encoded at different qualities.

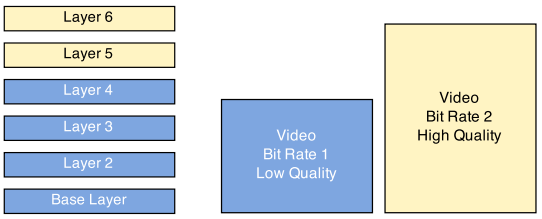

(a) Layers (SVC) (b) Versions (AVC)

Figure 2: Layers and versions for one time segment

A HAS application can also use Scalable Video Coding (SVC) instead of non-scalable coding. SVC encodes a video segment into a base layer and one or more enhancement layers using temporal, spatial, or quality scalability methods. The base layer represents a low-quality version of the video segment and every additional enhancement layer improves the quality of the segment. With a HAS application using SVC, the client can choose the bit rate of a video segment by the number of enhancement layers it requests, whereas with non-scalable video coding, the client chooses the bit rate by requesting a different file. Figures 2(a) and 2(b) show a video segment encoded in multiple layers and in multiple versions respectively. Decoding the lower four SVC layers would give a video quality comparable to decoding the smaller AVC video file, whereas decoding all six SVC layers would give a quality comparable to decoding the larger AVC video file.

A HAS client has to be aware of the available layers or versions on the server to determine what requests it can make. A media content description (MCD) or media presentation description is a file describing video segment information, such as timing, URL, video resolution, and bit rate. A HAS client can retrieve the MCD and use its information to make decisions on which layers or version to request. A widely supported standard that specifies the allowed content and structure for a MCD, but not which codec to use, is MPEG-DASH.

Deciding what bit rate to request for a certain video segment is one of the major challenges in HTTP Adaptive Streaming. When using non-scalable video coding, the version to retrieve is usually decided once per segment, since replacing an already retrieved version with a higher quality version renders the former obsolete and the bandwidth used for it wasted. With scalable video coding, a client generally has two options; the client can request additional layers for segments that will be played in the near-future, or it can request layers for segments more ahead of time [2]. The first option tries to maximize the current playback quality, at the cost of playback disruptions when large fluctuations in available bandwidth occur. The second option prioritizes uninterrupted playback, although it may result in a lower-than-necessary playback quality overall. Finding an optimal compromise in between these extremes is essential for a HAS application using SVC.

The authors of [1] show that two major commercial players (Smooth Streaming and Netflix) and one open-source player (Adobe's OSMF) make suboptimal decisions for the video bit rate they choose to retrieve under fluctuating bandwidth conditions. This exemplifies the difficulty in making optimal decisions for which version of a video segment should be retrieved in a HAS environment.

Scalable video coding has existed for more than a decade and it has been supported by the popular H.264/AVC standard since 2007 [10, 14]. Furthermore, there are implementations of HAS that use SVC, such as iDASH (improved dynamic adaptive streaming over HTTP), and they show that SVC allows for more optimal retrieval decisions than non-scalable HAS solutions [13]. Despite these benefits of SVC, video streaming providers such as Netflix currently use non-scalable video coding exclusively [8].

3 Efficiency and quality metrics

This section compares scalable and non-scalable video coding with each other on efficiency and quality metrics. Section 3.1 compares AVC and SVC on their coding complexity; Section 3.2 analyses the quality of experience obtained with scalable and non-scalable coding; Section 3.3 discusses the storage efficiency of both coding techniques; Section 3.4 discusses the caching performance of videos encoded as layers compared to videos encoded as versions; and Section 3.5 concludes this part by giving an overview of the comparison between SVC and AVC on the analysed metrics.

3.1 Coding complexity & efficiency

SVC decoders generally require 10 to 50% more computational resources than single-layer H.264/AVC decoders to obtain the same target resolution and bit rate, depending on the resolution ratios between the layers, the number of layers that are employed for interlayer prediction, and the bit rate [15]. This can be especially relevant for mobile devices for which energy efficiency is a key requirement.

Furthermore, since AVC is more widely used and better supported, it is more likely that hardware acceleration is available for it. Hardware acceleration allows the same or a better video quality while using less computational resources and less energy. This is again especially relevant for mobile devices.

3.2 Quality of experience

SVC can be incorporated in a live streaming environment [16]. Such a system provides a higher and more stable playback quality compared to AVC-based systems in high-delay networks with small buffer sizes [3, 12]. In a high-delay network with an application that has a small buffer size, video segments will be played back shortly after being retrieved. Therefore, there is generally insufficient time to retrieve a different version of a previously retrieved segment with non-scalable video coding. This results in the application choosing a quality lower than the available bit rate to avoid gaps in video playback when the available bandwidth decreases. With scalable video coding, the client application can refrain from requesting a higher layer when the available bandwidth abruptly decreases. Since this decision can be made without affecting previously retrieved layers, scalable video coding allows for a higher base quality in small-buffer/high-delay environments.

3.3 Storage efficiency

The advantage of a higher storage efficiency is that a higher number of different videos fit in a cache and that less bandwidth is required to transmit a video. With SVC, every set of consecutive layers that starts with the base layer can be decoded into a valid video, whereas with non-scalable video coding, only a complete segment can be decoded into a valid video. The partial independence between layers that scalable video coding provides, causes SVC to have a storage overhead compared to non-scalable video coding [16]. When achieving two-layer spatial or quality scalable encoding in SVC, SVC has an overhead below 10% compared to H.264/AVC-encoded files that receive similar quality of experience ratings [9, 14].

Thus, SVC has a small overhead per video compared to AVC if different-quality versions have to be available. However, the different versions are represented by different subsets of the layers, whereas non-scalable video coding requires a different video file for each version. Therefore, storing different-quality versions requires less storage capacity when using SVC than when using AVC.

3.4 Caching performance

In a caching environment, the origin server is the server that hosts all the original content, whereas the caching server usually hosts only part of the content. The caching server is located close to the end user and the bandwidth between the end user and the caching server is usually abundant. When a client requests a video from a caching server, the latency is low and the bandwidth is high compared to requesting the video from the origin server.

The closer location of caching servers decreases the chance of congestion, since the packets traverse fewer links in general and therefore also fewer links that could be congested. Furthermore, the original server is less likely to be the bottleneck since some requests can be answered by the caching servers.

Caches have limited space and are therefore usually unable to store all multimedia data available at the origin server. Choosing which content to store and in what format is key in obtaining a high cache hit ratio, since requests for files not present in the caching server are forwarded to the origin server. Smaller version-based files allow the storage of more files compared to larger layered files. However, having a base layer that is queried often improves caching performance of layers over versions [6].

Video streaming providers often provide videos in different qualities for better network performance, which can be achieved with different versions or multiple layers. The question with regards to caching performance is whether the highest cache hit ratio is obtained by caching versions, layers, or both in different situations.

In [13], the authors show the increase in cache hit ratio when SVC is used instead of AVC in a video-on-demand service that uses adaptive streaming. The authors of [6] researched whether caching layers or versions gives a better cache hit ratio, both when the request distribution for videos is known a priori and when it is not. They examined the throughput when using only versions, only layers, and when using two different mixed strategies using both layers and versions. The first mixed strategy caches a version the first time a video is requested. When a different quality of that video is requested, layers are used to answer the request. Those layers are cached and the previously cached version is removed from the cache as soon as it is no longer in use. The second mixed strategy streams a version when a video is requested the first time, but this version is not cached. The second time this video is requested, the corresponding layers are cached.

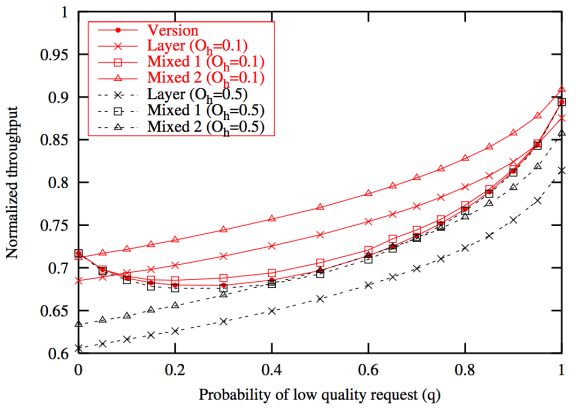

Figure 3: Adaptive caching for varying probability of low quality request (Figure from [6])

The authors observed that the second mixed strategy performed better than the first in general, because with the first strategy, versions sometimes remain in the cache because they still receive requests, preventing the more efficient layers to enter the cache. They observed that better cache hit ratios are obtained when caching layers instead of versions, except when a video is requested only once. Caching videos that will be requested only once, decreases the hit cache ratio. Therefore, videos requested for the first time with an unknown popularity, should be streamed directly from the origin server to the client. Streaming them as a version, requires less bandwidth compared to streaming layer(s) that provide an equal quality. Only when a video is requested a second time, should the layer(s) corresponding to that request be stored on the caching server. Such a mixed strategy using both versions and layers provides the best result in terms of caching efficiency.

Figure 3 shows one of the results of [6]. It shows the normalized throughput for the different strategies using only versions, only layers, and the two mixed strategies in relation to the probability of a request being for the low quality option of a video. The figure confirms that the second mixed strategy in general delivers the highest throughput, except when all requests are either for a low-quality or for a high-quality video, or when the layered encoding overhead compared to versions is too large.

3.5 Overview of efficiency and quality metrics

Table 1 shows the comparison of H.264/AVC with SVC on the properties discussed in the previous sections.

Property	Comparison
Coding complexity	AVC is 10-50% more efficient than SVC with regards to the computational resources use by the decoder
Hardware support	Many devices support AVC hardware acceleration; SVC hardware acceleration is often not supported
Quality of experience	SVC provides a better quality of experience in live-streaming environments with small buffers or a high latency
Storage efficiency	AVC is maximum 10% more efficient than SVC when comparing videos in a single quality
	SVC is more storage efficient when multiple qualities of the same video are stored
Caching performance	Caching only layers gives a better cache hit ratio than caching versions; versions could be used for direct streaming from the origin server

Table 1: Comparison of SVC and H.264/AVC

4 Operating cost

The operating cost for a video streaming service is important for the viability of the service and for the competitive advantage over competitors. These costs depend partially on the coding technique used for the provided videos, especially with regards to the required storage capacity and bandwidth usage.

The previous section explained that storing different qualities of a video requires less storage space when using SVC than when using AVC. However, since bandwidth is a significant cost factor for video streaming services, the small overhead that SVC has per bit stream over AVC outweighs the benefits of decreased storage requirements for popular videos. For often requested videos, SVC turns out to be more costly than AVC, whereas SVC is less costly for less-requested videos [7]. Therefore using SVC for less popular videos and AVC for more popular ones is an opportunity for cost reduction.

A side note here is that [7] does not take caching costs and caching performance into account. When using caching servers, the total amount of required storage space increases per caching server whereas the bandwidth cost increases per streamed video. Therefore the storage efficiency benefit of SVC increases with the number of caching servers used, whereas the bandwidth cost is independent thereof.

5 Conclusion

This paper outlined the differences between scalable and non-scalable video coding with a focus on the comparison of both with regards to multimedia streaming applications.

This paper shows that H.264/SVC decoders generally require 10 to 50% more computational resources than H.264/AVC decoders and that the AVC decoder is more likely to be hardware accelerated. Therefore, AVC provides a better coding efficiency in general.

SVC provides a more stable playback quality in live-streaming environments where the playback buffer is small. SVC also provides a better storage efficiency if multiple video qualities are required. For a single video quality, SVC has an overhead of maximum 10% compared to AVC.

This paper explained that the operating cost of a popular video streaming service is usually higher using SVC then when using AVC to encode the videos, because bandwidth costs are a more significant factor than storage costs. However, caching layers results in a better caching performance than caching versions, which influences the operating cost and reduces the chance on congestion at the origin server.

Based on the properties outlined in this paper, AVC provides a better fit if coding complexity needs to be minimized or when hardware support at the decoder is important. Moreover, sending a single video quality requires more bandwidth using SVC compared to AVC and since bandwidth is costly, the operating cost of a video-on-demand (VoD) service using AVC may be lower. However, when using caching servers, SVC provides a higher cache hit ratio, which leads to less congestion and a higher throughput for the service without the need for a better performing origin server. A VoD service using caches with layers provides a better quality of experience to its users than a VoD service using AVC with caching or a VoD service using no caching at all. Further research is needed to examine if the higher cache performance with SVC outweighs the additional bandwidth usage needed for SVC when considering the total cost of the service. If it does not, the service provider can weigh the better quality of experience of its users against the additional cost.

6 Future Work

Despite the advantages SVC has with regards to playback quality, caching efficiency, and storage efficiency as presented in this paper, AVC is still the most used video encoding technique. Decoding all existing AVC videos and recoding them with SVC is a resource-intensive process. Transcoding from AVC to SVC and vice versa could prove a less resource-intensive solution [4].

A downside to encoding a video into layers with SVC is that a missing layer causes all higher layers to be unusable. Multiple Descriptive Coding (MDC) is a coding technique using forward error correction to mitigate this problem. MDC has a significant overhead compared to SVC caused by the extra bytes needed for the forward error correction. Adaptive Layer Distribution is another coding technique which tries to minimize this downside by providing a midway between efficiency and error resilience [11].

HTTP adaptive streaming adapts the video quality and requested bit rate based on the playback buffer it maintains. Prolonged network congestion causes the playback buffer length to decrease which often causes the client to request lower bit rate segments from the server. Simultaneously, congestion affects TCP's congestion avoidance algorithm which may decrease the bandwidth the application can use. The interaction between TCP's congestion avoidance algorithm and adaptation algorithms used in HAS services are areas for future research.

References

[1] S. Akhshabi, S. Narayanaswamy, A. C. Begen, and C. Dovrolis. An experimental evaluation of rate-adaptive video players over HTTP. *Signal Processing: Image Communication*, 27(4):271–287, 2012.

[2] T. Andelin, V. Chetty, D. Harbaugh, S. Warnick, and D. Zappala. Quality selection for Dynamic Adaptive Streaming over HTTP with Scalable Video Coding. *Proceedings of the 3rd Multimedia Systems Conference on - MMSys '12*, page 149, 2012.

[3] N. Bouten, S. Latr, J. Famaey, and F. D. Turck. Minimizing the Impact of Delay on Live SVC-based HTTP Adaptive Streaming Services. *Ifip/Ieee Im 2013*, pages 1399–1404, 2013.

[4] J. D. Cock, S. Notebaert, P. Lambert, and R. V. D. Walle. Architectures for Fast Transcoding of H.264 AVC to Quality-Scalable SVC Streams .pdf. 11(7):1209–1224, 2009.

[5] J. Famaey, S. Latré, N. Bouten, W. de Meerssche, B. De Vleeschauwer, W. Van Leekwijck, and F. De Turck. On the merits of SVC-based HTTP Adaptive Streaming. In *Integrated Network Management (IM 2013), 2013 IFIP/IEEE International Symposium on*, pages 419–426. IEEE, 2013.

[6] F. Hartanto, J. Kangasharju, M. Reisslein, and K. Ross. Caching video objects: Layers vs versions? *Multimedia Tools and Applications*, 31:221–245, 2006.

[7] H. Kalva, V. Adzic, and B. Furht. Comparing MPEG AVC and SVC for adaptive HTTP streaming. *Digest of Technical Papers - IEEE International Conference on Consumer Electronics*, pages 158–159, 2012.

[8] J. Martin, Y. Fu, N. Wourms, and T. Shaw. Characterizing Netflix bandwidth consumption. *2013 IEEE 10th Consumer Communications and Networking Conference, CCNC 2013*, pages 230–235, 2013.

[9] T. Oelbaum, H. Schwarz, M. Wien, T. Wiegand, I. Communication, and T. U. Berlin. SUBJECTIVE PERFORMANCE EVALUATION OF THE SVC EXTENSION OF H . 264 / AVC Department of Electrical Engineering and Information Technology , TU M ÂÍ Institute of Communications Engineering , RWTH Aachen University , Aachen , Germany 3. *Image Processing*, pages 2772–2775, 2008.

[10] J.-R. Ohm. Advances in Scalable Video Coding. *Proceedings of the IEEE*, 93(1):42–56, 2005.

[11] J. J. Quinlan and A. H. Zahran. ALD : Adaptive Layer Distribution for Scalable Video. *MMSys 2013*, pages 202–213, 2013.

[12] Y. Sanchez, T. Schierl, C. Hellge, T. Wiegand, D. Hong, D. De Vleeschauwer, W. Van Leekwijck, and Y. Le Louédec. Efficient HTTP-based streaming using Scalable Video Coding. *Signal Processing: Image Communication*, 27(4):329–342, 2012.

[13] Y. Sánchez, T. Schierl, C. Hellge, T. Wiegand, and D. D. Vleeschauwer. iDASH : Improved Dynamic Adaptive Streaming over HTTP using Scalable Video Coding. *MMSys 2011*, pages 257–264, 2011.

[14] H. Schwarz, D. Marpe, and T. Wiegand. Overview of the Scalable Video Coding Extension of the H . 264 / AVC Standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 17(9):1103–1120, 2007.

[15] H. Schwarz and M. Wien. The scalable video coding extension of the h. 264/avc standard. *IEEE Signal Processing Magazine*, 25(2):135, 2008.

[16] M. Wien, H. Schwarz, and T. Oelbaum. Performance analysis of SVC. *IEEE Transactions on Circuits and Systems for Video Technology*, 17(9):1194–1203, 2007.

Biometric Identification Methods

Juho Saarela
Student number: 84425K
juho.saarela@aalto.fi

Abstract

Due to the problems inherent in password-based identification, biometrics, such as fingerprints or iris recognition, is becoming an increasingly common method of identifying and authenticating users in everyday contexts. This paper discusses some of the most common or promising methods of biometric identification available and assesses their usability and reliability.

Some common biometrics not discussed in this paper include hand geometry, voice recognition, DNA testing and skin spectroscopy.

KEYWORDS: biometric, identification, authentication

1 Introduction

In the modern world, a person is expected to remember many of strong passwords, all unique to the specific service in question. However, in practice people tend to use merely a few different passwords with only limited strength, as numerous studies show [19, 10, 18].

Biometric identification is able to counter this vulnerability, since it is impossible for a person to "forget" his biometric data because is an intrinsic property of the person in question. Ideally, biometric identification would not only free users of having to remember passwords, but also of the risk of having their the passwords stolen or cracked.

However, although it is possible to bypass a simple fingerprint check by having a copy of the print or by simply cutting off the actual finger of the user, the more advanced biometric identification methods typically prevent this kind of unauthorised usage. When authentication requires multiple biometric identifiers, forgery becomes extremely difficult.

However, losing biometric data to an outside entity is significantly more critical than losing a password. There are a limited number of biometric traits in humans (10 fingers, 2 eyes, etc.), and if those all are compromised, secure biometric authentication is no longer possible without a liveness detection setup [17]. Similarly, a user losing his biometric trait, for instance in an accident, could potentially lock a user out of a system.

Since biometric identification provides a means to potentially make authentication easier for the user as well as increase security, it is definitely a worthy path of research. However, it is not a technology without risk. Should a technology such as automated face recognition become sufficiently advanced, it would make maintaining privacy considerably harder in a heavily surveilled environment, such as most modern urban areas.

2 Biometric Identification in General

The basic idea in biometric identification is to **enroll** a user into the system by providing a **template**, an initial **sample** against which future samples can be compared [33]. For instance, in the case of the fingerprint registry of the police, the template is created when a person applies for a passport or commits a crime the first time, and subsequent samples are then matched against that template.

While in the registry of the police a person's fingerprint is associated with his other private information, such as social security number, this is not necessary in biometric identification. Using biometrics, anonymous identification is possible if no data associated with the biometric template reveals the identity of the user [32]. This allows the creation of user registries that are strictly personal but do not reveal the identities of the users. The difference between such a biometric system and a traditional password-based identification is that passwords can be shared while biometrics cannot.

Because of the numerous varying factors involved in taking samples, the various samples taken from a person are never exactly the same as the template [33]. Possible reasons for the differences are numerous. Overall, the combination of the human factors, varying sampling environments and potentially different equipment all have an effect on the sampling process and can thus lead to greatly varying sample quality [11]. Acceptance **threshold** is the maximum disparity between a template and a sample so that the system accepts a match [32].

Two of the most important metrics involved in biometric identification are **false rejection rate** (FRR) and **false acceptance rate** (FAR). False rejection (also known as type I error) occurs when a user has a template in the system, but the new sample does not match the template. False acceptance (also known as type II error) occurs when a sample is incorrectly accepted by the system as a match to a template. False rejection and false acceptance rates determine how often these events occur [33].

Equal error rate (EER) is found by getting the acceptance threshold where the FAR and FRR are equal, giving a relatively good indication of the overall robustness of the system. [1] However, since false acceptance is more detrimental to security than false rejection, the acceptance threshold should be stricter than what it is for the EER.

It should be noted that false acceptance rate only accounts for someone else using his own biometrics to try to access the system. In other words, even if a system has a 0 % false

acceptance rate, it might be trivial to spoof, making the system insecure [2].

While increasing the acceptance threshold decreases the false rejection rate, it also increases the false acceptance rate. Thus, it is important to find the optimal value to ensure that the system does not unnecessarily reject authorised users and at the same time still effectively prevents unauthorised access [33].

It is necessary to take great care with storing the biometric data. If the templates, for example fingerprints or iris images, are stored in a raw format, hacking into the database could potentially compromise all of the accounts linked to the stored fingerprints. To solve this issue, P. Li et al. [17] propose a method which would encode the fingerprint data using cryptography, similar to how passwords are typically hashed.

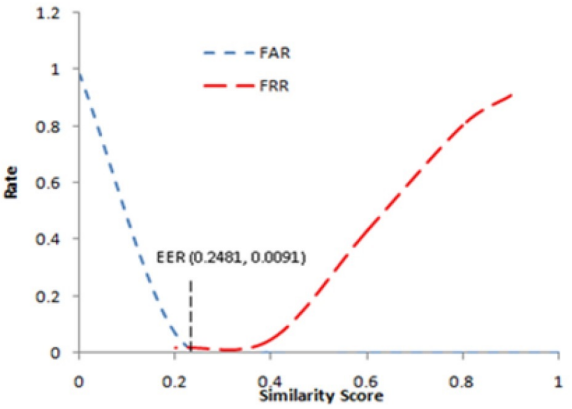

Figure 1: A graph showing the FAR and FRR of an ATM fingerprint system. The similarity score is the acceptance threshold used. [1]

3 Identification methods

This section introduces the various common biometric identification methods available and covers their basic features and operation. Fingerprints are the most commonly used biometric trait, but the others are also in use in contexts that require a higher level of security or where fingerprints are not applicable.

3.1 Fingerprints

Probably the most widely used form of biometric identification method is fingerprint identification. Since fingerprints are almost unique to an individual and it is relatively straightforward to obtain a sample, it has maintained its popularity despite other emerging identification methods. Some modern mobile phones have integrated fingerprint sensors for authentication [9]. Lee [16] estimates that the chance of two people having the same fingerprint is less than one in a billion.

A serious problem with fingerprint authentication, as Jan Krissler [14] has demonstrated, is that fingerprints can be duplicated and successfully used from smudges on a mobile phone screen, with instructions being available online for everyone to follow [13]. A public system using fingerprint authentication would be especially vulnerable, as a malicious user would be able to copy the fingerprints of previous users. In addition, it is also possible to produce fake fingerprints from high-resolution photographs.

There are various methods to counter fake fingerprints being employed to bypass authentication. The electrical impedance of a living finger can accurately distinguish it from synthetic fingers, and it is possible to integrate such a device into the actual fingerprint sensor without loss of imaging quality or the need for an external device [29]. Since fake fingers tend to be more rigid than actual fingers, pressure compared to the area of the fingerprint can also reveal whether the fingerprint was made using a fake finger. Other properties that can help determine whether a fingerprint is genuine or not include temperature, electrical conductivity, pulse, perspiration and odour [21].

Fake fingers can also be detected on the software level. It is possible to recognise 90 % of fingerprints made using fake fingers of various materials by simply analysing the image sample [7].

Another issue with fingerprints is that there are multitude of conditions that can make sampling difficult. Fingers can be wet or dry and they may be damaged, resulting in the samples produced to be of low quality. [22]

In public use, although people do not mind using fingerprints because of the ease of use and familiarity, the lack of hygiene is of some concern, as the user needs to physically touch the sensor. [24]

Overall, despite the many drawbacks, fingerprints are a very practical biometric trait. With careful use, fingerprints can be employed in many contexts, including ATMs, and can reach ERR of 0.91 %, as Fig. 1 shows [1].

3.2 Vein pattern matching

The veins in human hands and fingers are unique and thus can be used as biometric identifiers [22]. While fingerprints and iris images are relatively trivial to spoof without sophisticated liveness detection methods, the vein pattern of a hand or a finger of a person is very difficult to spoof.

The hand veins are always distinct, and even twins do not have matching vein patterns. Likewise, the vein patterns of left and right hand are different. In addition, a sample can only be taken from a live body and they do not change over time. Since the veins are not visible to human eyes, and considering how difficult it would be to replicate a synthetic copy of a human hand, it is a biometric trait that is almost impossible to steal. Since the vein imaging does not require a physical contact with any device, it allows hygienic public usage. [6] Unlike fingerprints, which are susceptible to degradation and damage, vein patterns are a more durable biometric identifier.

The benefit of vein pattern matching is that it can be easily combined with other biometrics, such as finger dorsal texture matching. Fig. 2 shows how both the finger vein pattern and the dorsal texture can be captured at the same time. The IR light passes through the finger, allowing imaging the veins,

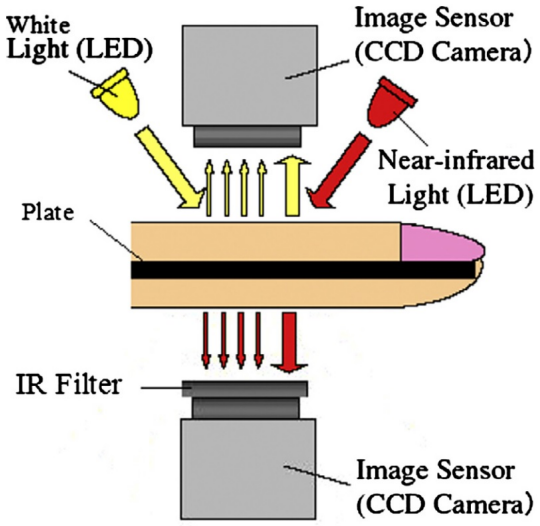

Figure 2: A figure of the setup where both the finger vein and dorsal texture patterns are captured. [34]

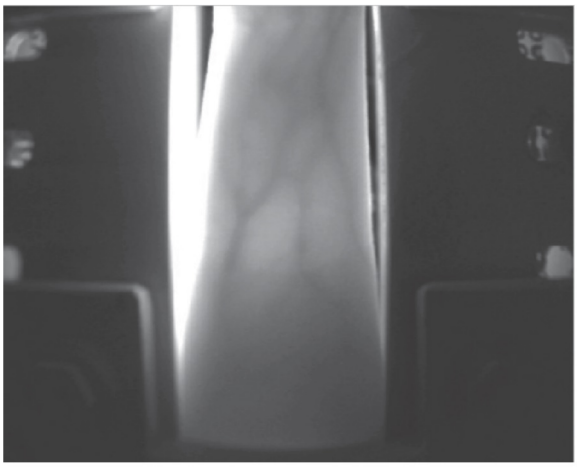

Figure 3: A finger vein image. [34]

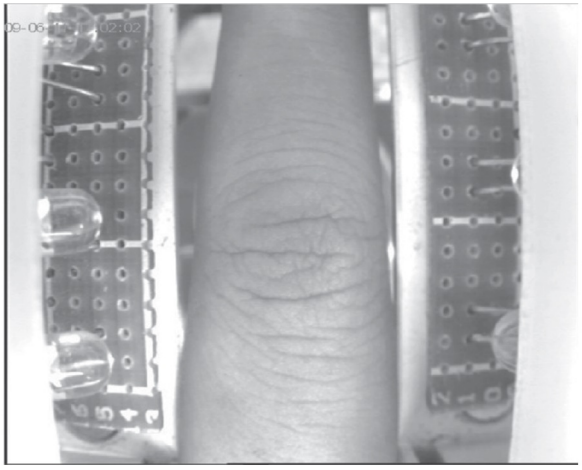

Figure 4: A finger dorsal texture image. [34]

while the white LED is used to capture the finger texture. Fig. 3 shows the finger vein image and Fig. 4 the dorsal texture image.

Vein pattern matching is highly accurate and difficult to spoof, thus offering a trustworthy option for biometric authentication. When finger vein pattern is combined with finger dorsal texture matching, EER as low as 0.435 % can be reached. [34] Even with only finger vein matching, an EER of 0.5 % is possible. [6] However, when imaging finger veins, the rotation of the finger must be constant to get accurate results, making it cumbersome and error-prone to use. [34]

3.3 Iris

The iris pattern is unique from eye to eye, making it a good biometric trait similar to fingerprints. The iris contains many details and thus provides much data for biometric analysis. They are also immutable, do not degrade over age and have a natural protection provided by the cornea. [23]

Typically iris images are captured using digital cameras and near-infrared lights to illuminate the eye. The near-infrared lights allow the camera to better capture dark, heavily pigmented irises while being unintrusive to the human eye. [4] Lu et al. [20] have demonstrated that mobile phones can also be used for capturing the processing the images.

Detecting spoofing can be performed at both hardware and software level. On the hardware level, the liveness of the iris can be verified, for instance, by observing the pupil reacting to light changes and by having multiple lights emitting various frequencies and simultaneously analysing the spectrum and 3D structure of the iris pattern. [9]

Software-based methods analyse the image and try to detect whether the sample was captured from a printed picture or an actual eye. This is achieved by signal processing, local binary patterns, or other various algorithmic methods. In a competition for the best liveness detection algorithm held in 2014, the winner group achieved an FAR of 0 % and an FRR of 0.5 % with the test set used in the competition. [28]

Generally, hardware-based liveness detection performs better than software-based approaches, but also costs more and is less flexible. [9]

Imaging irises poses many issues, because various obstructing factors, such as reflections, eyeglasses, eyelashes, image noise and look direction, may degrade the sample quality. Wang et al. [31] have managed to create a system that reaches an EER of 1.83 % using a challenging iris image database.

When using specialised equipment that captures the iris patterns of both eyes and has an integrated facial recognition system, iris pattern matching can reach an EER of 0.131. The same system when using only one iris as the biometric achieves an EER of 3.29 %, and 0.36 % when using both irises. [12] This shows that multimodal biometric authenti-

cation offers significant improvements over using just a single biometric trait.

3.4 Retina

The retina is a light-sensitive layer of tissue located at the back of the eye and its use as a biometric is based on the uniqueness of its vein pattern. Aside from DNA, the vein pattern of the retina is the most reliable and stable biometric trait and varies even between identical twins. [25]

Although retinal scans are used in areas requiring high security, such as military or government complexes, the need for expensive equipment combined with the cumbersome sampling process prevents widespread use. [25] The vein pattern of a retina can achieve an EER of 0 % [15].

3.5 Facial features

Using facial features as a biometric trait presents many challenges not inherent in most other biometric traits. The background, illumination and facial expression can be different from sample to sample, on top of which aging, varying hairstyles and items such as glasses can lead to discrepancies. Although the human brain is able to recognise thousands of faces, machine-based recognition is still facing significant challenges. [26]

Machine-based face recognition can be based on many different algorithms. A simple approach is to match either the entire face or the individual features of the face. However, this requires a lot of memory and computational performance and it is difficult to extract specific facial features. [26] The more advanced methods use concepts such as eigenfaces[26] or multiscale local phase quantization[5].

Though it is possible to spoof a facial recognition system with images or displays, there are methods to algorithmically efficiently spot fakes from live samples with high accuracy. Some of these methods are based on pure image analysis while others take the overall movements of the face and facial features into account[8]. Some face detection systems use 3D data, increasing accuracy and making spoofing more difficult[3].

Since no recent research seems to provide a common metric EER for a face recognition system, it is difficult to compare its overall performance to other biometrics. However, using different face image databases, each containing images of various quality, Wagner et al. [30] have managed to achieve a recognition rate of at least 95.1 %. When using a database of faces with sunglasses obstructing the view, the rate drops to 40.9 %. In a survey conducted in 2006, some facial recognition systems achieved an EER of 3 %. [3]

3.6 touch-based identification

It is possible to identify users according to their way of using a mobile phone's touch screen. The pressure, duration and touching width of the presses are different for everyone and can be analysed to determine the identity of the user. Faking another user's patterns is very difficult, because the attacker would need to have similar neurology, finger softness and using habits. [27]

Biometric trait	EER
Fingerprint	0.91 % [1]
Finger vein pattern	0.5 % [6]
Finger vein pattern + Finger dorsal texture	0.435 % [34]
Single iris	1.83 % [31]
Both irises	0.36 % [12]
Both irises + face	0.131 % [12]
Facial recognition	3 % [3]
Retina	0 % [15]
Touch-based	N/A

Table 1: A table showing the equal error rate of various biometric authentication systems. No research papers of touch-based identification provided EER values.

In a system proposed by Seo [27] et al. the system initially monitors user's behaviour and collects data of input patterns and sends it to a server. A server is used because intense calculation is required to analyse the patterns and learn the typical characteristics of the user's way of entering input. If the user's patterns differ too much from the template on the server, the user needs to authenticate using an additional method, such as a password or a fingerprint.

A huge benefit of touch-based identification is that it requires no additional effort from the user. While fingerprint or iris scan need an additional action from the user, touch-based authentication can be active at all times without obstructing usability in any way. Furthermore, because there is no need for additional hardware, it is also cheap to employ. The system can identify users with an accuracy of n̈early 100 %̇ [27]

4 Data Analysis

Comparing various biometric authentication systems is difficult, because the EER metric seen in Table 1 does not give a comprehensive picture of the overall quality of the system. Hypothetically, there could be two systems that have an equal EER of 2.5 %. One of these could be configured to use an acceptance threshold that gives the system an FAR of 0 % and an FRR of 3.0 %. Meanwhile the other might not be able to achieve an FAR of 0 % without the FRR increasing to, say, 80 %. It is clear that in a case like this, the latter system is significantly worse. Due to the inconsistent way many papers presented their results, it was difficult to obtain anything except the EER to compare them, however. A more descriptive metric would be listing the FRR at various fixed FAR values, for example at 1, 0.1, 0.01 and 0.001 %.

In addition, the use contexts of various biometric authentication systems can be vastly different. For example, capturing iris images with a mobile phone is bound to produce different results from using specialised equipment for the task.

The comparison does not reveal practicality of the systems, either. Even though retina-based biometric authentication is superior to the other biometric authentication methods listed here, it requires highly specialised equipment and, therefore, is unsuitable for widespread usage, as fingerprints are, for example.

Using multimodal authentication provides better results. A system should check more than just one biometric trait whenever it is feasible. For example, a mobile phone could be used to check both the fingerprint and the iris of the user, but implementing a hand or finger vein matching system in mobile context would be impractical.

5 Conclusion

New technologies concerning biometric identification combined with the widespread use of mobile phones and other devices able to capture samples open up new avenues concerning how people will authenticate themselves in the future. Already more and more people use fingerprints instead of passwords to unlock their devices and it is only a matter of time until passwords become the exception instead of the rule.

It is important to note, however, that the error rates of various biometric identification methods are based on the sometimes quite limited amount of test data used in the research and thus they should be viewed with caution. Even so, the results are very promising and already provide a great number of biometric traits and their combinations that can be used to reliably identify a person.

References

[1] I. Babatunde and A. Charles. A fingerprint-based authentication framework for ATM machines. *J Comput Eng Inf Technol 2: 3. doi: http://dx. doi. org/10.4172/2324*, 9307:2, 2013.

[2] B. Biggio, Z. Akhtar, G. Fumera, G. L. Marcialis, and F. Roli. Security evaluation of biometric authentication systems under real spoofing attacks. *IET biometrics*, 1(1):11–24, 2012.

[3] K. W. Bowyer, K. Chang, and P. Flynn. A survey of approaches and challenges in 3d and multi-modal 3d+ 2d face recognition. *Computer vision and image understanding*, 101(1):1–15, 2006.

[4] K. W. Bowyer, K. Hollingsworth, and P. J. Flynn. Image understanding for iris biometrics: A survey. *Computer vision and image understanding*, 110(2):281–307, 2008.

[5] C. H. Chan, M. A. Tahir, J. Kittler, and M. Pietikainen. Multiscale local phase quantization for robust component-based face recognition using kernel fusion of multiple descriptors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(5):1164–1177, 2013.

[6] A. Chandra et al. Finger vein based user identification using differential box counting. *IJRCCT*, 3(1):138–142, 2014.

[7] J. Galbally, F. Alonso-Fernandez, J. Fierrez, and J. Ortega-Garcia. A high performance fingerprint liveness detection method based on quality related features. *Future Generation Computer Systems*, 28(1):311–321, 2012.

[8] J. Galbally, S. Marcel, and J. Fierrez. Image quality assessment for fake biometric detection: Application to iris, fingerprint, and face recognition. *IEEE Transactions on Image Processing*, 23(2):710–724, 2014.

[9] D. Gragnaniello, C. Sansone, and L. Verdoliva. Iris liveness detection for mobile devices based on local descriptors. *Pattern Recognition Letters*, 2014.

[10] S. Haque, M. Wright, and S. Scielzo. A study of user password strategy for multiple accounts. In *Proceedings of the third ACM conference on Data and application security and privacy*, pages 173–176. ACM, 2013.

[11] A. K. Jain, A. Ross, and S. Prabhakar. An introduction to biometric recognition. *IEEE Transactions on Circuits and Systems for Video Technology*, 14(1):4–20, 2004.

[12] Y. G. Kim, K. Y. Shin, E. C. Lee, and K. R. Park. Multimodal biometric system based on the recognition of face and both irises. *Int J Adv Robotic Sy*, 9(65), 2012.

[13] J. Krissler and Chaos Computer Club. Chaos computer club breaks apple touchid. `http://www.ccc.de/en/updates/2013/ccc-breaks-apple-touchid`, 2013. Accessed: 2015-16-04.

[14] J. Krissler and Chaos Computer Club. Ich sehe, also bin ich ... du. `http://media.ccc.de/browse/congress/2014/31c3_-_6450_-_de_-_saal_1_-_201412272030_-_ich_sehe_also_bin_ich_du_-_starbug.html`, 2014. Accessed: 2015-16-04.

[15] S. M. Lajevardi, A. Arakala, S. A. Davis, and K. J. Horadam. Retina verification system based on biometric graph matching. *IEEE Transactions on Image Processing*, 22(9):3625–3635, 2013.

[16] S. W. Lee and B. H. Nam. Fingerprint recognition using wavelet transform and probabilistic neural network. In *International Joint Conference on Neural Networks, 1999. IJCNN'99.*, volume 5, pages 3276–3279. IEEE, 1999.

[17] P. Li, X. Yang, H. Qiao, K. Cao, E. Liu, and J. Tian. An effective biometric cryptosystem combining fingerprints with error correction codes. *Expert Systems with Applications*, 39(7):6562–6574, 2012.

[18] Z. Li, W. Han, and W. Xu. A large-scale empirical analysis of chinese web passwords. In *Proc. USENIX Security*, pages 1–16, 2014.

[19] Z. Liu, Y. Hong, and D. Pi. A large-scale study of web password habits of chinese network users. *Journal of Software*, 9(2):293–297, 2014.

[20] H. Lu, C. R. Chatwin, and R. C. Young. *Iris recognition on low computational power mobile devices*. INTECH Open Access Publisher, 2011.

[21] E. Marasco and A. Ross. A survey on anti-spoofing schemes for fingerprints. *ACM Computing Surveys*, 47:1–36, 2014.

[22] C. Nandini, C. Ashwini, M. Aparna, N. Ramani, P. Kini, and K. Sheeba. Biometric authentication by dorsal hand vein pattern. *International Journal of Engineering and Technology*, 2(15):837–840, 2012.

[23] M. Negin, T. A. Chmielewski Jr, M. Salganicoff, U. von Seelen, P. Venetainer, and G. G. Zhang. An iris biometric system for public and personal use. *Computer*, 33(2):70–75, 2000.

[24] G. Peevers, R. Williams, G. Douglas, and M. A. Jack. Usability study of fingerprint and palmvein biometric technologies at the ATM. *International Journal of Technology and Human Interaction (IJTHI)*, 9(1):78–95, 2013.

[25] S. Qamber, Z. Waheed, and M. U. Akram. Personal identification system based on vascular pattern of human retina. In *Biomedical Engineering Conference (CIBEC), 2012 Cairo International*, pages 64–67. IEEE, 2012.

[26] S. Ravi and S. Nayeem. A study on face recognition technique based on eigenface. *International Journal of Applied Information Systems*, 5(4), 2013.

[27] H. Seo, E. Kim, and H. K. Kim. A novel biometric identification based on a users input pattern analysis for intelligent mobile devices. *International Journal of Advanced Robotic Systems*, 9, 2012.

[28] A. F. Sequeira, H. P. Oliveira, J. C. Monteiro, J. P. Monteiro, and J. S. Cardoso. Mobilive 2014-mobile iris liveness detection competition. In *Biometrics (IJCB), 2014 IEEE International Joint Conference on*, pages 1–6. IEEE, 2014.

[29] T. Shimamura, H. Morimura, N. Shimoyama, T. Sakata, S. Shigematsu, K. Machida, and M. Nakanishi. Impedance-sensing circuit techniques for integration of a fraud detection function into a capacitive fingerprint sensor. *Sensors Journal, IEEE*, 12(5):1393–1401, 2012.

[30] A. Wagner, J. Wright, A. Ganesh, Z. Zhou, H. Mobahi, and Y. Ma. Toward a practical face recognition system: Robust alignment and illumination by sparse representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(2):372–386, 2012.

[31] N. Wang, Q. Li, A. A. A. El-Latif, T. Zhang, and X. Niu. Toward accurate localization and high recognition performance for noisy iris images. *Multimedia tools and applications*, 71(3):1411–1430, 2014.

[32] J. Wayman, A. Jain, D. Maltoni, and D. Maio. An introduction to biometric authentication systems. In *Biometric Systems*, pages 1–20. Springer, 2005.

[33] A. C. Weaver. Biometric authentication. *Computer*, 39(2):96–97, 2006.

[34] W. Yang, X. Huang, F. Zhou, and Q. Liao. Comparative competitive coding for personal identification by using finger vein and finger dorsal texture fusion. *Information Sciences*, 268:20–32, 2014.

Survey of ARM TrustZone applications

Pawel Sarbinowski
Student number: 466819
pawel.sarbinowski@aalto.fi

Abstract

ARM TrustZone security extensions were introduced almost a decade ago in the ARMv6 architecture and since then they have received a lot of attention as a security primitive in low-cost hardware. The main concept introduced in TrustZone is the separation of the CPU operating mode into a "secure" and "non-secure" mode. The two modes are isolated via hardware-based access control features without the need for a separate co-processor for the secure environment. Commonly, TrustZone technology is leveraged to run minimal pieces of software to which sensitive operations are delegates, while the regular operating system and applications are run in the non-secure mode.

Demands for trusted computing in commodity computing platforms, including handsets, tablets, wearable devices and even embedded systems, have stimulated the industry research (e.g.[1],[2]). Innovative applications are also being explored increasingly by academia. Some of those include payment protection technology, digital rights management, BYOD (Bring Your Own Device) security enforcement, and a plethora of other security solutions. This paper surveys current applications of ARM TrustZone technology in both industry and academia. It also explores the ongoing standardization efforts involved in the domain, and identifies potential open research problems in the area.

KEYWORDS: ARM TrustZone, Trusted Execution, Integrity Monitoring, Linux, Virtualization, Mobile Security

1 Introduction and Motivation

In the past two decades the usage of mobile and embedded devices across all categories of computer ecosystems has increased rapidly. These devices, used by consumers or as infrastructure elements, are usually interconnected and create or contain vast amounts of information. At the same time, various stakeholders are striving to provide better security in the mobile ecosystem, motivated by several requirements. [3] [4] Some of those include:

- regulatory requirements, e.g. that the radio frequency parameters defined during manufacture are stored securely and remain unaltered. Recently, a bill in the state of California was passed that forces all mobile phones that are being sold to have a "kill switch" mechanism that will be able to shutdown and destroy the phones remotely in the event of theft or loss [5].

- standardization requirements such as the mandate that the International Mobile Equipment Identifier (IMEI) remains unchanged. This is also important for uniquely identifying devices and dealing with thefts.

- business requirements such as protection of digital content using Digital rights management (DRM) that is backed up by hardware security mechanisms or enforcement of phone operator subsidy locks to ensure that subsidized phones given to subscribers as part of a contract cannot be used with other mobile operators. Also remote attestation of the integrity of a client or employee computer might be necessary in cases of remote collaboration. This would ensure that the computing base used on the remote devices can be trusted and that the boot system and operating system have not been tampered with.

- users security requirements that present the need for secure storage of various credentials and other private information (e.g. contacts) on mobile devices not only for login credentials but also for applications such as secure payments. Another scenario of interest can be performing integrity checks for anti-virus and anti-theft software to subsequently notify the user if something is altered.

- developers interest in securing certain parts of their code that hold sensitive information (e.g. keys for remote API).

These raised security concerns have led to various system-level solutions employing hardware to run sensitive code or store private information in Trusted Execution Environments (TEEs). TEEs utilize specialized hardware to provide a secure, integrity protected environment for processing and storing information.

Currently, the TEE's functionality would not typically be available to third-parties, other than the Device manufacturers. OEMs take advantage of the TEE to protect resources (e.g. licensing metadata in the case of DRM) and functionality, such as crypto operations, from an adversary that has administrative control of the normal Operating System. However, there are many use cases where third-party application, and not just the OEM, would also benefit from security features provided by a TEE. Examples of such applications are One-time Password generation for two-factor authentication, User Authentication mechanisms, Secure Mobile transactions for e-payments and even enforcement of Access Control for applications containing private data such as e-health apps. Alternative hardware architectures for implementing these kind of scenarios typically utilize a trusted

co-processor (Trusted Platform Module or TPM) or hardware module (e.g. a SIM card) but a separate co-processor also leads to higher system costs and in some cases to a loss or difficulty in programmability.

ARM TrustZone security extensions have been available for some time already, since the ARM11 Core which was released in 2002. They are available in the large majority of new mobile and embedded devices with ARM processors. TrustZone's main advantage, compared to other solutions that use a separate on-chip or off-chip co-processor, is that it is inexpensive (i.e. no extra hardware chips are needed) and it takes advantage of the full processing power of the CPUs and not just a small subset of a slower secure chip. Both proprietary and open source TEE OS's and TEE applications are being built on top of TrustZone. At the moment, there are ongoing efforts (noted later in 6) for easing the process of developing and debugging on top of TrustZone. This paper explores some of the implementations utilizing TrustZone from both academia and industry.

The structure of this survey is as follows: Section 2 provides an overview of TrustZone technology, presents the key concepts of the architecture and discusses some of the challenges in the field. Sections 3 and 4 present some existing relevant applications utilizing TrustZone and then 5 and 6 discuss about related works and efforts to standardize the TrustZone API available for third party developers. Finally, Section 7 provides some concluding remarks.

2 Background

Existing hardware security solutions can be separated to two groups: those that utilize separate trusted co-processors in a different chip either outside or inside the CPU die and those that utilize single-chip solutions where the hardware security is incorporated into the CPU architecture itself and there is no second co-processor. Secure co-processor architectures can be split furthermore based on the hardware they use to: Smartcards used in systems such as ATMs, SmartTVs and communication equipment(e.g. SIM cards); Trusted Platform Modules (TPM) which have a separate secure co-processor usually soldered on a PC board and are the primary means of providing standardized trusted computing features to PCs; and Hardware Security Modules (HSM) that may have more than one cryptoprocessors inside and offer multiple levels of security not only against software attacks but also against physical attacks.

Some of the double-chip CPU architectures include: Intel TXT, which uses a TPM and measures the integrity of software and platform components (i.e. code, data structures, configuration files and more) in order to allow management applications to make trust decisions correspondingly, AMD PSP (Platform Security Processor), which uses a dedicated co-processor that features ARM TrustZone technology to guarantee features like secure boot, trusted execution and more.

Examples of some platforms having a single chip are: AEGIS [6], that utilizes memory integrity verification, encryption/decryption of off-chip memory and a secure context manager; XOM [7] (eXecute Only Memory), where the processor is equipped with a private key, known only to the processor manufacturer, and the corresponding public key is published. Hence, anyone can encrypt the software by using the public key but only the processor having the corresponding private key can decrypt and execute the software. All software is stored encrypted in memory when it is executed; Intel SGX (Software Guard Extensions)[8], which constitutes a set of new CPU instructions that can be used by applications to specify protected regions of code and data referred to as enclaves. An enclave is a protected area in the application's address space, which provides confidentiality and assures integrity even in the presence of privileged malware. Accesses to the enclave memory area from any software not resident in the enclave are prevented; and finally ARMs' TrustZone platform.

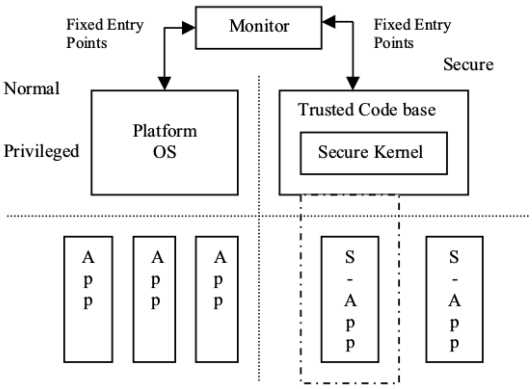

Figure 1: TrustZone architecture[9]

The ARM TrustZone architecture (Fig. 1) serves as a much more flexible single-chip alternative, with a fully functioning CPU whose secure component is freely programmable, compared to the double-chip solutions that have a fixed programming interface and usually feature set for secure applications to use. Programming a TEE on top of a double-chip TPM or HSM is also feasible but there are drawbacks in terms of performance since the secure co-processor is a more limited chip due to its cost. TrustZone introduces the concept of two logical processing modes for the ARM CPU, referred to as "secure world" and "normal world" mode. Each mode is isolated via hardware-based access control features from the other one. Typically, TrustZone technology is leveraged to run small, security-specialized portions of code in the secure world mode, whereas the conventional operating system and applications are run in the normal world mode. The distinction between the worlds is implemented in hardware thus prevailing software protection rings between user-level and kernel-level code are not affected. The operating system does not have to be aware of the two cpu modes for the applications to utilize them. The hardware separation is also propagated over the system bus to peripheral devices, memory controllers and cpu registers. Thus, when secure mode is active, the software running on the CPU has a view of the whole system that is isolated from the software that is running in non-secure mode.

The main benefit of this approach is the minimal loss in processing power of the secure cpu component (which is

now the same cpu) and the extensible functionality a programmer can implement on the cpu instead of the hardwired actions previously found on TPMs. To transition between the two worlds a special *Secure Configuration Register* (SCR) is used. Specifically, the Non-Secure bit in the register (SCR.NS) is set to 1 to transition from the secure to the non-secure world. Code running in normal mode cannot directly change the NS bit, therefore to enter to the secure mode again, a set of calls is defined and includes secure monitor calls, interrupts and external memory system aborts. Secure monitor calls are handled in the secure world and act as a sort of cross-world system call.

For completeness we will also define some additional terms such as: Rich Execution Environment (REE) is the full-feature operating system such as Linux, Windows, OS X, Android or iOS; TEE OS is the operating system running in the secure world and is usually comprised of a microkernel, a memory management unit and code implementing an API for Client/Trusted apps; TEE is a superset of that and essentially includes all the functionality (in hardware and software) that is necessary for the controlled isolation of secure apps from the REE; Trusted Application (TA) is an application that is running inside a TEE and provides certain functionality to Client Applications; finally Client Application (CA) is an application running in the normal world (REE) that accesses a TA's functionality or offloads a security critical part of it's own to a running TA.

The main features of the TrustZone platform can be split into: Boot Integrity, Secure Storage, Device Identification, Isolated Execution, Device Authentication and an API to access the TZ features. A quick overview of these components follows:

Boot Integrity - Acts as a building component for the security of the entire system. Verifies the integrity of the software at each stage of the boot (bios firmware, bootloader, operating system). The boot code's signature hash is checked against a securely stored hash that was signed by the device manufacturer. The device manufacturer's public key is used to decrypt the signature and to compare the two hashes. The comparison of the hashes continues for each subsequent boot stage and if they at some point differ, the device stops the boot (in the case of secure boot) or stores the different measurements for later usage by the OS or applications (in the case of authenticated boot).

Secure Storage - Uses cryptography to preserve the confidentiality and integrity of application data. Usually the encrypted data is stored in other insecure peripherals (e.g. a hard disk) but it is encrypted with a key that is stored inside the secure non-volatile memory of the CPU. All the cryptographic methods are also executed in the context of the secure world.

Device Identification - Uses the secure storage feature to verify and preserve multiple unique identifiers for each device. A base identity is defined and certified by the device manufacturer, which stores it in the secure non-volatile memory. This base identity can then be used to verify and sign new assigned identities that can be used for e.g. e-ticketing applications.

Isolated Execution - Runs trusted applications inside the secure world separated from each other and from the normal world. It also guarantees that any code and data running in the secure world is protected at run-time from access even from privileged code running in the normal operating system. A minimal TEE Operating system is used to manage Trusted applications running in the secure world as well as communication with them. Trusted applications that intend to run in the secure environment will have to be signed by the device manufacturer.

Device Authentication - Verifies the integrity of the system to third parties that need attestation. In this case, the system's state (O.S. running, boot software signature and trusted application) is signed with a device specific key known to the device manufacturer before-hand. The signed device authentication can then be verified by a third party by using the public key of the manufacturer that signs a pre-defined set of known system configurations (OS, boot code, apps etc) that are acceptable for the device, thus proving the integrity of the system.

API - The TEE operating system, running in the secure world, provides an programming interface for communication between Client Applications (CAs) in the "insecure" OS and Trusted Applications in the TEE, and even between Trusted Applications in the TEE. In the latter case, the isolation is provided by privileged mode code part of the TEE OS, not distinct hardware features.

2.1 TEE implementations

The ARM TrustZone extensions [10] define mainly the hardware architecture necessary for a secure computing component. To develop an application on top of those extensions a TEE is necessary. Although ARM provides some basic examples [11] of how to utilize the TZ security monitor to develop simple applications on top of TrustZone, in the long-term a more complex dedicated TEE OS running in the secure world will be a more robust solution capable of managing several Trusted Applications at the same time. Several proprietary TEE Operating Systems are already built on top of the TrustZone architecture to provide the necessary functionality and management of Trusted Applications.

Mobicore OS was initially developed by a company called Giesecke & Devrient (G&D), which together with the ARM Secure Services Division and Trusted Logic Mobility (TLM) established Trustonic [12]. Now, Trustonic provides both a (GP compliant) TEE operating system and a TEE Directory. TEE Directory is a framework for remote deployment and management of Trusted Applications built on top of Trustonic's TEE. It adheres to the GP remote management protocol standard. Trustonic provides an SDK for its TEE to developers (after registration) and acts as a Trusted Service Manager for the Trusted Apps that are deployed.

Qualcomm has also implemented a proprietary TEE OS of it's own called Qualcomm Secure Execution Environment (QSEE) [13] on top of TrustZone as provided in it's S4 Snapdragon 8XXX processor series. Besides the TEE, Qualcomm's environment also provides several proprietary Trusted Applications such as: StudioAccess for DRM in media streamed by partners among which are Amazon, Hulu, Netflix and more; Enterprise and BYOD for securing corporate data on employees devices with device-unique crypto

keys; SafeSwitch that acts as a remote kill switch in case of theft; Authentication by using biometrics based on Fast IDentity Online (FIDO) standard. Finally QSEE provides a TA that acts as a hardware back-end for Android's Keystore API.

Finally Solacia has also implemented SecuriTEE OS [14] as a proprietary TEE that adheres to the GlobalPlatform standards and is mostly used for secure payments and content protection.

All the aforementioned implementations are closed source, require licensing fees to use or develop on them and thus restrict the development of TEE Applications by more programmers.

3 Applications utilizing TrustZone

Existing applications based purely on TrustZone-provided features (i.e. trusted applications built on top of the TEE OS) are described in the following papers[[9],[15],[16],[17],[18]].

Motivated by the increase in both mobile device usage and in the ease of digital content distribution in the past decades Hussin. et. al.([9]), examines the usage of TrustZone for secure storage and enforcing of DRM policies. It utilizes an existing DRM distribution model found in Symbian OS. In this model, however, DRM metadata files and vital application data (such as media files, game levels, data controlling the application logic) are stored in the TrustZone-based secure storage. Furthermore, the DRM License Manager, which is in charge of regulating application execution through licenses and protecting application data, is run inside the secure world to ensure integrity of execution. The authors also propose an E-Pass application based on this DRM model, where the ticket issuer and the user will share a signed electronic pass. The pass will be stored transparently in the user's secure hardware and will resist tampering.

A platform, backed by TrustZone, for anonymous electronic payments is presented by Picker and Slamanig ([16]) where they try to unify the convenience of electronic payments with the anonymity of cash. Traditionally in electronic payments, a user first purchases electronic credit via some sort of e-banking and then he/she uses the credit on the provider's services (e.g. in the context of a transport payment system). To avoid the linking of users' banking accounts to their transport information the authors introduce signed electronic tokens that a user can buy anonymously with cash. The tokens are in the form of printed QR-codes that a user can scan with his smartphone. The signed token is then stored in the TZ secure storage to preserve privacy even in the event of the phone loss. The transport provider will also have to keep a corresponding list of the tokens issued for usage but there is no need to keep any type of personal information for the users, thus preserving their private information. The sensitive computations and storage would run as a Trusted Application (Trustlet) in the context of TrustZone.

Rijswijk-Deij. et. al([18]) - investigate the possibility of using a device having TrustZone capabilities to implement One Time Passwords (OTP) for two-factor authentication with a comparable level of security. To allow the comparison, a conceptual model of user interaction with OTP applications is introduced and studied. The model assumes there is a trusted path to the display, all user input passes through a TrustZone secure peripheral bus and that the Memory Management Unit support protected memory separation between the secure and the normal world. They show that a trusted path can exist between user input and display of OTP credentials, albeit with two disadvantages. First, the chip manufacturer is the external entity that must be trusted so that the system can work, and, second, that TZ does not come with a dedicated software implementation by itself, complicating its usage by developers.

4 REE OS Hardening

Research applications that effectively build another abstraction layer on top of the TEE OS are presented below. They aim mostly at creating a middle layer that either secures certain procedures transparently or provides a much easier to use API for applications to do so.

TrustZone-based Real-time Kernel Protection (TZ-RKP) [1] is an approach aiming to provide enhanced security to the Operating System kernel by transferring its security monitor to the TrustZone secure world. Consequently the security monitor that handles all access control is protected from attacks originating in the normal world where the rest of the kernel is running. TZ-RKP transfers the control over privileged system functions to the secure world where each access request is inspected before being granted. Examples of such malicious actions can be requesting for system sensitive data, hiding malicious processes or escalating the privileges of a malicious app. Existing hypervisor solutions achieve the same goal but are themselves also prone to attacks from processes running in the normal world. The target kernel runs in normal world while TZ-RKP, containing the security monitor, runs in the secure world. By running the security critical part of the kernel in the hardware backed secure world, TZ-RKP achieves a compromise between security and usability and since it is already deployed in Samsung Galaxy devises it is a tested and feasible approach. To guide the flow of code execution through the access monitor and memory management unit, residing in the secure world, TZ-RKP utilizes hooks in the kernel. This provides full control over the memory management of the normal-world kernel as well as live interception and monitoring of critical events that can be denied if they negatively impact the security of the system. To further ensure that the kernel hooks used are not altered or bypassed by a malicious process, TZ-RKP also: maps the memory containing the kernel code as read-only; defines the memory area containing dynamically allocated kernel data (such as virtual-to-physical memory translation tables) as read/write only by the kernel code and as non-executable; allows only trusted applications that are signed, verified and inspected to run in the TEE. Extensive testing confirmed TZ-RKP's effectiveness and benchmarks performed showed that the overhead of this system is minimal and ranges between 0.2% and 7%.

A security enhancing framework is introduced by Zu. et. al.([15]) designed for embedded systems running Linux. This development is motivated by the increase in

usage of embedded systems in pervasive computer environments (e.g. networking stacks, Internet of Things devices, instruments for automation and devices used in aeronautic research). The proposed framework includes a Linux Security Module that hooks on the Linux kernel and enforces additional access policies based on security models such as Bell-La Padula and the Domain and Type Enforcement model. To ensure that the security policies cannot be tampered with without authorization, the framework is mainly implemented as a Trusted application running in the secure world of TrustZone and provides an API for the LSM hooks of the kernel. This acts as an extra layer of Mandatory Access Control similar to the way modules like SELinux, AppArmor and Smack work but since it is backed on the TZ hardware, it is significantly harder for attackers to bypass without physical access.

A virtualization platform design to secure embedded and mobile systems is proposed by Johaness Winter ([19]). The core idea of it is to split insecure software into isolated user-space VMs that communicate with secure-world trusted engines. These are managed by a Linux based kernel running in the secure world that enforces MAC (mandatory access control) and isolated execution with SELinux policies. Each of these Trusted engines has a properly defined interface for communicating with other trusted engines. The isolation of the Trusted Engines is based on the features TrustZone offers i.e. secure peripherals cannot be accessed by non-secure software and non-secure software cannot access secure memory without authorization from the secure-world. The proposed architecture handles all the low-level details, such as dispatching secure monitor calls, and enforces restrictions on the resource usage of the non-secure mode software code. The platform uses a secure boot loader that authenticates the secure world linux kernel image and measures it. This measurement is then compares to a Reference Integrity Metric (RIM) certificate that is attached to the kernel image. Only if the RIM certificate matches to the integrity measurements taken, will the boot continue and control will be handed to the Operating system running in the non-secure world. To communicate with secure code any secure monitor call that is invoked by non-secure code is handled by the user-space VM supervisor. This VM supervisor acts as a middle layer between non-secure apps and secure-world Linux kernel, thus minimizing the amount of processing the secure kernel has to do. Eventually the authors plan to integrate the TrustZone based virtualization framework with the KVM virtualization framework found in Linux kernels.

A Trusted Language Runtime (TLR) is proposed, implemented and evaluated by Santos. et. al. ([2]). The TLR aims to act as a mechanism for isolating security-sensitive application logic from the rest of the application, the Operating system as well as other applications. It is based on a lighter .NET Micro runtime Framework (NETMF) designed for embedded devices, thus managing to keep the trusted computing base code reasonably small. The framework handles low level details some of which are secure memory or cpu resource sharing, interrupt handling across domains and coordination of the execution flow of the apps between secure and non-secure worlds. By obscuring these details the framework allows developers to easily split their security-sensitive

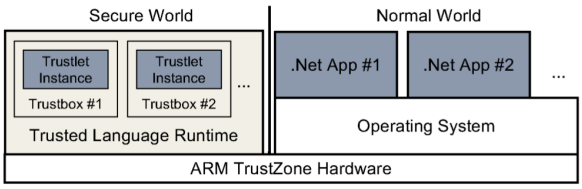

Figure 2: TLR high-level architecture[19]

code into classes that transparently run in a trusted environment, isolated from the rest of the app, the operating system and from other Trusted applications (by using TrustZone features) and in a development framework (.NET) which they are already familiar with. Figure 2 provides a better overview of the framework.

To maintain a small TCB certain restrictions were introduced to the runtime. Code that is running on the Trusted Instance of the runtime (in the secure world) can only perform computations or encrypt/decrypt data associated with trusted classes. It cannot access peripherals. Otherwise the TLR would have to include drivers to access them which would increase the codebase significantly. Sample use cases of apps utilizing the TLR secure features include One-time Password generators, User authentication and Secure Mobile Transactions (both based on a challenge-response protocol running on trustlets) and enforcement of access control to secure data used or collected by e.g. e-health applications.

5 Related Work

Because very little work had been done on open source software for TrustZone systems Johannes Winter ([17]) explores system-level development on cost-efficient arm devices, particularly in a classroom environment. The goal is to utilize this knowledge to teach TEE programming in classrooms. Obtaining such knowledge had been troublesome since most TrustZone enabled board manufacturers either do not provide any technical support or they do only after signing a non-disclosure agreement which complicates matters significantly for academics. In the paper the authors: analyze the TrustZone platform of a specific ARM board, discuss the problems that arose while developing a micro-kernel prototype that would run in the secure world, present methods to test the basic TrustZone functionality of a board as well as a simple architecture for initial development of apps running in normal mode, and, finally, demonstrate the development of a secure boot-code while also analyzing the board's internal boot ROM structure.

Motivated by the lack of tools for debugging and developing TEE applications, McGillion. et. al. ([20]) introduce Open-TEE, a virtual TEE implemented entirely in software that emulates the TEE behavior as defined in the GP [21] standards. By utilizing a software-based TEE emulator developers can avoid the cost of proprietary TEE software development kits or expensive hardware debugging tools, while at the same time taking advantage of already known, reliable debugging tools (such as GDB or LLDB) and de-

velopment environments. Thus, after fully debugging and testing a Trusted Application, they will be able to compile it to any GP-compliant hardware TEE. To make it as simple as possible to deploy and work with, the framework was designed as a set of processes each implementing specific TEE operations (e.g. Trusted application, TA manager, TA launcher), thus the time and resource overhead to develop with Open-TEE is minimal. A small scale, but extensive, user study was also performed and showed positive results on the ease of use of Open-TEE and its benefit for the TEE development life-cycle. Open-TEE is freely available as open-source software under the Apache-V2 license.

Dan Rosenberg in his recent work covered a newly found vulnerability in the QSEE TEE software [22] which affects "all known Android devices that support TrustZone and utilize a Qualcomm Snapdragon SoC, with the exception of the Samsung Galaxy S5 and HTC One M8, which have been patched". The vulnerability is based on a bug in the bounds-checking the Secure Monitor performs on calls made by client apps. It allows a kernel-level application to write data to arbitrary secure-world memory, thus completely compromising any trusted application running in the TEE. Similar vulnerability research was also done in Azimuth Security in 2013 to exploit QSEE and bypass the Secure Bootloader in Motorola Devices.

6 Standardization Efforts

The restrictions of proprietary TEE implementations2.1 and the difficulty in programming and debugging in them led to an increased need for both standardization of the TEE API and for open source implementations of TEE Operating Systems and developer tools. Steps already taken in that direction include the following.

"GlobalPlatform is a cross industry (Visa, MasterCard, NTT, ARM and others), non-profit association which identifies, develops and publishes specifications that promote the secure and inter-operable deployment and management of multiple applications on secure chip technology." [21]. Around 2011, GP published specifications for the TEE client API (used by client applications running in the normal world), the TEE System Architecture and TEE Internal API (used for internal communication between the corresponding TEE parts running in secure and non-secure world). "GlobalPlatform has launched a TEE compliance program. This offers assurances to application and software developers and hardware manufacturers that a TEE product will perform in line with the GlobalPlatform standards and as intended. It also promotes market stability by providing a long-term, inter-operable and industry agreed framework that will evolve with technical requirements over time". Several TEEs including Mobicore OS, QSEE, SecuriTEE, OP-TEE and SierraTEE are already partially if not full GP compliant and therefore applications built on those specifications should be able to inter-operate between all platforms.

Besides the GP standardization effort, there have also been endeavors to create open source TEE operating systems as an alternative to the proprietary solutions mention in section 2.1.

Linaro [23] is an open organization focused on improving the security of Linux on ARM. Linaro along with STMicroelectronic developed OP-TEE, which is an Open Source TEE. It is based on a proprietary TEE solution that was Global Platform certified on some products in the past thus compliance with the GP standards is quite extended. Currently the secure TEE code runs on 32 bit only but there are plans to add support for 64 bit and to also upstream the OP-TEE kernel driver to the Linux kernel.

SierraTEE and SierraVisor [24] are an open source TEE and hypervisor respectively that offer a comprehensive solution from secure boot to application management and are maintained by the Sierraware. They support applications in C, C++ and Java and are easy to integrate with android and other mobile platforms. They also have the option to use dedicated cores for normal and secure world OR share worlds between cores. It is a minimal TEE and offers protection against side channel attacks, supports several interrupt models, fast performance architecture, fast context switching and supports asynchronous IPC. It currently supports ARM11, Cortex-A9, and Cortex-A15 processors and is also available for 64bits.

7 Conclusion

In this report we examined the research and product development done by academia and industry, for the ARM TrustZone architecture. Overall although the initial motivations for the research (i.e. securing the interests of all the stakeholders) remains we can observe a shift in how this research is performed. From mostly proprietary and not openly documented products to more and more solutions that are based on the GP standard and are even open source. The research area around TrustZone is positively active. And the reason is that even though TrustZone was partially marketed as a "silver bullet" for mobile security and as a secure DRM platform it still encounters some difficulties. There is a shortage of open documentation about both hardware and software TZ related products which is mostly important for academia since proprietary products cannot be easily referenced in scientific work. Trusted application developers faced several issues both in terms of compatibility among the different hardware TEEs and in terms of costs to buy proprietary software development kits and debugging tools. This state is slowly changing for the best with open source efforts including: Open-TEE, OP-TEE, SierraTEE and more. The spread of open source solutions will definitely benefit the smaller developers that wish to utilize and benefit from the platform as well as smaller businesses or academia that do not have the resources to afford proprietary solutions. Finally, it should be noted that since TEEs and their APIs are still software solutions they can also be vulnerable to attacks (albeit on a different level) by malicious parties residing in the normal-world OS as was seen in 5. Therefore TrustZone should not be considered as a cure-all solution but as another extra layer of security. Extensive peer review of the TEE code and well-tested open source solutions would be vital in the long-term.

References

[1] Ahmed M Azab, Peng Ning, Jitesh Shah, Quan Chen, Rohan Bhutkar, Guruprasad Ganesh, Jia Ma, and Wenbo Shen. Hypervision across worlds: Real-time kernel protection from the arm trustzone secure world. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 90–102. ACM, 2014.

[2] Nuno Santos, Himanshu Raj, Stefan Saroiu, and Alec Wolman. Using ARM TrustZone to build a trusted language runtime for mobile applications. In *Proceedings of the 19th international conference on Architectural support for programming languages and operating systems*, pages 67–80. ACM, 2014.

[3] J.-E. Ekberg, K. Kostiainen, and N. Asokan. The Untapped Potential of Trusted Execution Environments on Mobile Devices. *Security Privacy, IEEE*, 12(4):29–37, July 2014.

[4] Atul Verma. Get Into the Zone: Building Secure Systems with ARM TrustZone Technology . Unpublished, Texas Instruments, February 2013. `http://www.ti.com/lit/wp/spry228/spry228.pdf`.

[5] Senate Bill No. 962 (kill switch bill). `http://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201320140SB962`. Accessed: 2015-01-30.

[6] Edward Suh, Dwaine Clarke, Blaise Gassend, Marten Van Dijk, and Srinivas Devadas. The AEGIS processor architecture for tamperevident and tamper resistant processing. *Massachusetts Institute of Technology*, 2003.

[7] David Lie, Chandramohan Thekkath, Mark Mitchell, Patrick Lincoln, Dan Boneh, John Mitchell, and Mark Horowitz. Architectural support for copy and tamper resistant software. *ACM SIGPLAN Notices*, 35(11):168–177, 2000.

[8] Frank McKeen, Ilya Alexandrovich, Alex Berenzon, Carlos V Rozas, Hisham Shafi, Vedvyas Shanbhogue, and Uday R Savagaonkar. Innovative instructions and software model for isolated execution. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy*, pages 1–1. ACM, 2013.

[9] WH Hussin, Reuben Edwards, and Paul Coulton. E-pass using DRM in Symbian v8 OS and TrustZone: Securing vital data on mobile devices. In *Mobile Business, 2006. ICMB'06. International Conference on*, pages 14–14. IEEE, 2006.

[10] ARM TrustZone whitepaper. `http://infocenter.arm.com/help/topic/com.arm.doc.prd29-genc-009492c/PRD29-GENC-009492C_trustzone_security_whitepaper.pdf`. Accessed: 2015-01-30.

[11] ARM TrustZone example. `http://infocenter.arm.com/help/index.jsp?topic=/com.arm.doc.faqs/ka15417.html`. Accessed: 2015-01-30.

[12] Trustonic - TrustZone and TEE. `https://www.trustonic.com/about-us/who-we-are`. Accessed: 2015-04-10.

[13] Qualcomm - QSEE. `https://www.qualcomm.com/products/snapdragon/security`. Accessed: 2015-04-10.

[14] Solacia - Securi-TEE. `http://www.sola-cia.com/en/securiTee/product.asp`. Accessed: 2015-04-10.

[15] Zu Yan-ling, Pan Wei, and Zhang Xin-guo. Design and implementation of secure embedded systems based on TrustZone. In *Embedded Software and Systems, 2008. ICESS'08. International Conference on*, pages 136–141. IEEE, 2008.

[16] Martin Pirker and Daniel Slamanig. A framework for privacy-preserving mobile payment on security enhanced ARM TrustZone platforms. In *Trust, Security and Privacy in Computing and Communications (TrustCom), 2012 IEEE 11th International Conference on*, pages 1155–1160. IEEE, 2012.

[17] Johannes Winter. Experimenting with ARM TrustZone. Or: How I Met Friendly Piece of Trusted Hardware. In *Trust, Security and Privacy in Computing and Communications (TrustCom), 2012 IEEE 11th International Conference on*, pages 1161–1166. IEEE, 2012.

[18] Roland Rijswijk-Deij and Erik Poll. Using Trusted Execution Environments in Two-factor Authentication: comparing approaches. 2013.

[19] Johannes Winter. Trusted computing building blocks for embedded linux-based ARM TrustZone platforms. In *Proceedings of the 3rd ACM workshop on Scalable trusted computing*, pages 21–30. ACM, 2008.

[20] OpenTEE. `https://github.com/Open-TEE`. Accessed: 2015-01-30.

[21] Global Platform TEE Specifications. `http://www.globalplatform.org/specificationsdevice.asp`. Accessed: 2015-01-30.

[22] Dan Rosenberg. QSEE TrustZone Kernel Integer Overflow Vulnerability. `https://www.blackhat.com/docs/us-14/materials/us-14-Rosenberg-Reflections-On-Trusting-TrustZone-WP.pdf`, jul 2014.

[23] STMicroelectronics and Linaro's OP-TEE. `https://wiki.linaro.org/WorkingGroups/Security/OP-TEE`. Accessed: 2015-01-30.

[24] Open Virtualization's SierraVisor and SierraTEE. `http://www.openvirtualization.org/`. Accessed: 2015-01-30.

Secure Public Instant Messaging: A survey

Dawin Schmidt
Student number: 466864
Dawin.Schmidt@aalto.fi

Abstract

Most of the public Instant Messaging (IM) services claim to provide end-to-end encrypted private chat. However, many IM services rely on the TLS protocol which does not protect against eavesdropping on the server side. Also, IM services usually do not safeguard against impersonation, allowing an attacker to fake IM identities.

Following the Snowden revelations, there has been a rapid growth of secure IM services that try to achieve private IM conversations. However, until now there has been no survey of secure IM services. This seminar paper surveys different IM protocols and applications and provides an analysis based on their security and privacy features.

KEYWORDS: Secure Instant Messaging, Privacy, Perfect Forward Secrecy, Deniability

1 Introduction

Instant Messaging (IM) is one of the most widely used online and mobile communication services. However, many current IM applications do not protect user's conversations. A decade ago, security did not play a major role in instant messaging and private communications could be easily sniffed during network transit. The most popular public domain IM services in those days, such as AOL Instant Messenger (AIM) or MSN Messenger, did not provide any security features. Nowadays, most of the IM solutions support Transport Layer Security (TLS) in order to provide strong client-to-server encryption and to protect against passive network attacks. However, TLS is not sufficient to protect instant messaging against eavesdropping because a malicious IM provider can still intercept the connection while relaying the messages from one user to another.

Even more embarrassing, intelligence services and governments are spying on anyone's encrypted communications. The Snowden revelations have confirmed that intelligence services such as the NSA and GCHQ are wiretapping IM [18] even though IM traffic is encrypted via TLS. In fact, the NSA is able to successfully bypass cryptography in order to spy on IM communication. For instance in 2008, the NSA collected up to 60.000 online sessions of Yahoo's Web-Messenger per day and massively gathered its buddy lists, which sometimes contain offline messages waiting to be delivered [18]. In addition, agencies also force big internet companies such as Skype to allow access to their internal networks in order to intercept IM network traffic [48]. Because of its weak security architecture, Skype further allows for possible governmental eavesdropping [50] as well as TLS man-in-the-middle (MITM) attacks performed by specialized network appliances [2]. Moreover, other technology companies such as Google [29] or BlackBerry [51] have been asked by governmental agencies to provide data about citizens' private communications or to add backdoors to their products [42].

According to Snowden's documents, the NSA runs a programme called "Bullrun", in which the agency aims to store huge amounts of encrypted internet traffic for later decryption [44]. The agency's goal is to decrypt such data in the future, either when new technologies such as quantum computing might be able to break encryption much faster or when the NSA manages to obtain the private encryption key. Insecure TLS cipher-suites may also be vulnerable to such actions if the key is compromised.

Intelligence agencies are not the only eavesdropper on IM communications: Internet companies also eavesdrop by either providing surveillance data to governments or by analysing their customers' communications for marketing purposes. The technology company WhatsApp, which is now owned by Facebook, is one example. Even though the firm recently announced they would support end-to-end encryption [25] in their famous IM application, the privacy of a WhatsApp conversation can no longer be fully trusted. Although the content of the messages are encrypted, WhatsApp can learn about the meta data of a conversation. In addition to the implicit knowledge of who is talking to whom, from where and for how long, WhatsApp is also able to gather information about users' actions and the actual length and language of the message by simply observing the sizes of the encrypted packets [11]. This kind of traffic analysis shows that encryption alone cannot protect privacy. Therefore, we need IM services which not only rely on encryption but make use of additional cryptographic primitives to protect user's privacy.

Following the Snowden revelations, there has been a rapid growth of secure IM applications recently. Additionally, several new secure IM protocol have been published. However, until now there has been no survey of secure IM protocols and applications. This seminar paper examines current secure IM protocols and services, and makes a survey based on the privacy and security features that these techniques offer.

This work does not discuss enterprise IM technologies because they have other requirements [52] and face different security threats [46]. In addition, the security analysis of group chat and file transfer mechanisms are beyond the scope of this seminar paper. The main focus remains on public do-

main IM protocols and applications for two-party communications.

This survey paper is organized as follows: Section 2 gives a technical overview of (secure) instant messaging and its related aspects. It also discusses major problems and challenges that are associated with basic IM technologies. Section 3 discusses secure instant messaging protocols and evaluates their different security properties. Five different solutions are presented: Off-the-record messaging protocol, TextSecure v2 protocol, SCIMP, Cryptocat and Bleep. Whereas section 4 reviews related work, section 5 discusses the main results. Finally, section 6 provides some concluding remarks.

2 Background

This section provides some preliminary material for the reader in order to understand the technical content that follows in section 3.

2.1 (Insecure) Instant Messaging

Instant messaging is a type of online-chat for exchanging text messages over a network and can be used either in a corporate or individual setting. In contrast to e-mail, IM is designed for a real-time communication and typically transmits text as it is being typed[1]. In addition, IM can be used for offline chat, placing phone/video calls, and transferring files. In case of offline-chat, the IM application either stores the offline messages and sends them once it is back online, or the IM server stores messages and sends them as soon as the IM application is connected to the network.

There are two communication architectures that IM solutions are modeled on: client-to-server and client-to-client (peer-to-peer). In a client-to-server communication model, an IM server manages the availability of the clients and relays messages between them. For authentication, each client shares a secret (e.g. a password) with the IM service provider. In addition, the server provides necessary information, such as the end-points' IP addresses, in case two clients communicate with each other in a peer-to-peer fashion (e.g. for audio/video chat or file transfer). There are also fully decentralized peer-to-peer IM networks available in which a central server is not necessary (e.g. Bleep).

Most of the IM protocols which exist today are proprietary and incompatible to each other, e.g. a WhatsApp user cannot talk to Skype users. Therefore, open-source IM protocols evolved to bridge this gap by acting as a gateway to other proprietary protocols. The main open-source instant messaging protocols are SIP for instant messaging and Presence Leveraging Extensions (SIMPLE)[2] and Extensible Messaging and Presence Protocol (XMPP)[3], which are standardized by the Internet Engineering Task Force (IETF). XMPP is a widely used protocol and supports transports[4], file transfers behind NATs[5], multi-user chat, serverless messaging[6], media sessions[7] and many other features.

Current instant messaging applications have basic built-in security and privacy features [32]. The following are some of the most important techniques:

Anti-spam protection. This mechanism prevents the user from receiving spam messages. For example, most applications require explicit user consent before adding a new contact to its buddy list and provide an option to blacklist malicious accounts. Furthermore, IM service providers typically block automated account creation which could be otherwise exploited to send spam messages.

Protection against malware spreading. Malware is typically transferred as files via instant messaging. In order to protect against malware spreading, IM applications require explicit user consent for accepting files. IM applications also notify users when new software updates are available. The updates usually consist of critical security bug fixes which could otherwise be exploited by malware. Similarly, antivirus software can also detect and block malicious IM file transfers.

Privacy settings. Most IM applications allow users to change privacy-related configuration settings. For example, IM applications allow an option to delete the chat history either automatically after a certain amount of time or manually. Moreover, IM applications require explicit user consent to publish private information such as online status or location information. The IM application may also require the user's login password before changing privacy-related settings in order to avoid unauthorized changes.

Authentication. Typically, IM applications use a password-based authentication mechanism to authenticate users. These user credentials are commonly encrypted via TLS during transit. Additionally, some IM applications support the OAuth 2.0 protocol or the Simple-Authentication-and-Security-Layer (SASL[8]) framework for authentication.

Encryption. Most of the IM services today support TLS encryption that can be enforced by the IM application or by the IM server. One example is the XMPP protocol, which uses TLS with a STARTTLS extension to protect the data stream.

2.2 Security and Privacy Threats to IM

Even though instant messaging faces many security and privacy threats [33] [32] [28], the default security mecha-

[1] http://www.realtimetext.org/rtt_in_detail/standards

[2] https://tools.ietf.org/html/rfc6914

[3] http://xmpp.org/about-xmpp/technology-overview/

[4] A XMPP-Transport is a gateway to other IM protocols in order to allow chat with other IM networks.

[5] http://xmpp.org/extensions/xep-0096.html

[6] http://www.xmpp.org/extensions/xep-0174.html

[7] http://xmpp.org/extensions/xep-0166.html

[8] https://tools.ietf.org/html/rfc2222

nisms of popular IM applications are not sufficient to protect against them. The main challenges are listed below:

Insecure connections. A major threat to instant messaging are insecure connections, which allow an adversary to trace private conversations. This is often caused by IM protocols which do not support encryption by default, e.g. the SIMPLE protocol does not consider any security mechanisms. Nevertheless, most IM protocols implement encryption by leveraging TLS. However, TLS may protect IM chat only to a certain extent. One drawback of TLS is that it does not guarantee end-to-end encryption when it is used in a client-to-server IM architecture. This is due to the fact that a TLS connection is only established between the client and server. In an IM conversation between two clients, one of the clients might actually have an unencrypted connection to the relaying server, enabling an attacker to sniff clear-text messages. If one communicating client is connected to another IM service provider, the server-to-server communication link might be also unencrypted, leading to the same passive attack vectors. Moreover, TLS does not protect against active attacks. For example, an attacker could perform man-in-the-middle attacks using a compromised IM server certificate[9].

Insecure endpoints. Although client-to-server IM communications are protected by TLS, the vulnerability might be at the endpoint. For example, one of the user's devices may be infected by malware, allowing an attacker to read chat messages.

Even though computationally expensive, some IM applications rely on asymmetric cryptography to encrypt messages. IM applications, which allow an attacker to relate a user's IP address to his/her public key, expose an additional attack vector. If an attacker manages to guess such relation, he/she could try to break into a user's remote computer to steal the private key. Some IM application do not hide such sensitive information, e.g. the user's presence status can leak his/her public key and automatic file downloads can reveal the corresponding IP address.

IM applications can also suffer from insecure default settings, e.g. by not encrypting messages by default. As a matter of fact, many user-convenient IM features are counterproductive to security. For example, most IM applications log old messages to create a message archive or store offline messages to be send once the destination is connected to the network. Another example are higher-level IM transport protocols which may leak user's metadata such as nicknames.

Impersonation. By pretending to be another person, criminals or malicious users can be a threat to a legitimate user. Because most of the IM server providers do not verify user identities, anyone can create accounts using any name[10]. By performing social engineering, a malicious person could easily create a personalized authorization request to ask the victim to add him/her to the victim's contact list, pretending to be a friend. Once the victim has added the malicious person, he/she could send a malicious hyperlink leading to a phishing website or transfer an infected file (which the victim would accept). The malware could then use the victim's buddy list as a propagation vector by spreading to all other contacts. Impersonation could be also misused to send spam messaging (SPIM[11]).

Malicious IM servers. Malicious IM servers are another major threat to instant messaging because they can store a copy of the session key and collect chat messages. There are two types of malicious IM servers: A compromised legitimate IM server and a rogue IM server which was set up by an attacker. By performing DNS spoofing attacks, an attacker could force clients to connect to his/her rogue IM server. If the IM application does not validate the server's certificate (e.g. by not checking the certificate against a DNSSEC secured DNS server via DANE or by certificate and public key pinning[12]), the attacker is able to eavesdrop on all communications.

Another security threat are proprietary IM servers because they lack public code audits and may contain serious security flaws or backdoors. Additionally, one needs to trust the company that maintains the IM server software. For example, companies typically control the key exchange infrastructure which makes it easy for governments to install MITM decryption keys (e.g. via a national security letter). Similarly, public XMPP server providers may also be a security and privacy threat as many of them store metadata. Many XMPP servers store the XMPP ID, the user's contact list, offline messages, IP addresses, the last user login, the total user online time, the amount of send/received packets and much more sensitive metadata. Older versions of a popular XMPP server application even stored the users' passwords in clear-text[13]. Good advice for choosing a secure XMPP server provider is therefore to read through its privacy policy and to check the security features of the server[14].

Proprietary IM applications. Proprietary IM applications may use non-standard cryptographic algorithms and may contain bugs. These non-standard techniques and bugs may be exploitable to leak user's privacy. Additionally, governments or intelligence agencies can force companies to add backdoors to their IM products [42].

Some of the proprietary IM applications are as follows:

- **WhatsApp** implements the open-source TextSecure protocol, yet the IM application is closed-source.
- **Skype's** TLS-based client-to-server architecture does not provide end-to-end encryption, which enables eavesdropping by design [48].
- **iMessage** has several weaknesses, thus allowing Apple to read encrypted IM chat [21] [10].

[9] https://en.wikipedia.org/wiki/Certificate_authority#CA_compromise

[10] On the Internet, nobody knows you're a dog: https://en.wikipedia.org/wiki/On_the_Internet,_nobody_knows_you're_a_dog

[11] https://en.wikipedia.org/wiki/Messaging_spam

[12] https://www.owasp.org/index.php/Certificate_and_Public_Key_Pinning

[13] https://www.ejabberd.im/plaintext-passwords-db

[14] https://xmpp.net/

- **SnapChat** had a security flaw, which allowed to bypass its self-destruct feature [5].

Denial of Service. Denial of Service (DoS) is a serious threat to instant messaging as it blocks clients from communicating with each other. Some IM protocols have security flaws, making them vulnerable for DoS attacks. Furthermore, an IM server could be a major target for distributed DoS attacks, affecting an entire IM network. Attackers could also perform DoS attacks against single clients by sending a large amount of messages from compromised accounts.

Unsuitability of PGP in IM Most IM applications only support data encryption from the client to the server via TLS. In order to guarantee end-to-end encryption, PGP encryption can be applied. Although providing strong confidentiality, PGP does not meet the requirements for secure IM communication and is therefore not recommended [31] [24] [40]. For instance, PGP does not support perfect forward secrecy, even though there have been attempts to add this feature to the official Internet standard [9]. This would allow an adversary to decrypt all previous messages once he or she obtains the private encryption/decryption key. Another issue with PGP is that it conflates non-repudiation and authentication because of the use of digital signatures. As a result, Bob can prove to a third-party that Alice told him a certain fact. Finally, new secure IM protocols are replacing PGP as an IM encryption technique. For example, the popular PGP IM application Psi has not been updated since 2012. Additionally, the official OpenPGP extension (XEP-0027[15]) for the XMPP protocol is obsolete, which can be seen as another sign that PGP in IM will be longer used in future. However, there are some use-cases for PGP in instant messaging, as for example in secure messaging systems such as Pond[16]. By using PGP encryption, users avoid to only rely on Pond's safety [31]. As a side note, this seminar paper will not consider S/MIME for secure instant messaging because it relies on a broken Certificate Authority (CA) model.

2.3 Requirements for secure IM

As discussed in the previous section, the basic security mechanisms of most IM solutions do not adequately preserve the user's security and privacy. In fact, a private and secure IM chat must have a similar security level as a face-to-face conversation between persons:

Confidentiality. Confidentiality guarantees that nobody is able to listen to Alice's and Bob's private conversation. Encrypting the communication channel ensures confidentiality. Additionally, IM protocols must provide end-to-end encryption. Thus, the IM server must not be able to learn about the encryption key which the endpoints use to secure the channel.

Perfect Forward Secrecy. Perfect forward secrecy (PFS) [27] ensures that it is impossible for someone else to find out what Alice and Bob talked about after their conversation has happened. Typically, this is achieved by short-term encryption/decryption keys which are generated only when they are needed and thrown away after usage. Another requirement must be that it is impossible to derive those keys from any long-lived key material.

Authentication. Another important property of private communications is authentication. Alice must be sure that she is really talking to Bob and that Bob is not pretending to be someone else. Typically, digital signatures and message authentication codes are used for message authentication. For entity authentication there are two common mechanisms in IM: Either by running the Socialist Millionaire's protocol [30]) or by providing a user-interface to verify public key fingerprints.

Deniability. Deniability allows a sender to send an authenticated message to a receiver. After message arrival, the receiver cannot prove to a third-party that such a message was sent by the sender [12]. For example, Alice might assume that her friend Bob is an FBI informant[17] who secretly keeps tracks of all her messages [23]. Therefore, Bob should not be able to prove to a third-party (e.g. to a court of law) that Alice sent any particular messages. If Alice and Bob use symmetric cryptography, both of them share the same private key for message authenticity. Thus, Bob could also create fake messages in Alice's name, making it impossible for him to provide valid proof. This concept is known as "weak" deniability or repudiability [7]. "Strong" (plausible) deniability or forgeability [7] is a related property, which assumes that not only Bob but everyone who listens to the communication channel could create fake messages. As a result, Alice can deny having sent a certain message because anyone else could have created messages in her name.

On the contrary, there is non-repudiation which is the opposite of deniability. For example, the third-party plugin for the popular IM application Pidgin, Pidgin-Encryption[18], uses digital signatures which are clearly checkable by a third-party.

Anonymity Anonymity makes sure that nobody else must notice a personal face-to-face conversation between Alice and Bob. Anonymous instant messaging is about hiding metadata, which could be otherwise easily correlated with other records to identify individuals [43]. In many cases metadata is sufficient to either convict [4] or even kill individuals (e.g. by U.S. drone strikes [47]). Anonymity can be achieved by decentralized architectures such as peer-to-peer networks or by onion routing technologies such as Tor.

3 Secure Instant Messaging

We choose five different IM technologies and discuss them into more detail regarding their privacy and security features: Off-The-Record messaging protocol (OTR), TextSecure v2

[15] http://www.xmpp.org/extensions/xep-0027.html
[16] https://pond.imperialviolet.org
[17] https://en.wikipedia.org/wiki/Hector_Xavier_Monsegur
[18] http://pidgin-encrypt.sourceforge.net/

protocol, Silent Circle Instant Messaging Protocol (SCIMP), Cryptocat, and Bleep.

3.1 Off-The-Record Messaging Protocol

The Off-The-Record (OTR) Messaging Protocol was introduced in [7] by Borisov, Goldberg and Brewer and has found wide adoption among a large number of IM applications. There are IM applications for standard PC as well as for mobile operating systems: Jitsi[19] (Windows, Linux, Mac OS X), Pidgin[20] (Windows, Linux), Adium[21] (Mac OS X), and ChatSecure[22] (Android, iOS).

OTR, whose latest version is 3.4[23], runs on top of IM transport protocols such as XMPP, ICQ, MSN, Yahoo, IRC[24], PSYC and many more. The protocol provides secure IM chat and encrypted file transfer, but does not yet facilitates multi-user group chat. OTR features most of the requirements of a private face-to-face conversation:

Confidentiality. Messages are encrypted with a session key using AES in counter mode (CTR). The communication link is end-to-end encrypted because the key exchange occurs between the two communicating clients only.

Perfect Forward Secrecy. The short-term session keys are negotiated using a Diffie-Hellman key exchange and are discarded after use. OTR does not utilize any long-term secret keys during session key generation.

Authentication. OTR uses a hybrid authentication approach by using digital signatures and message authentication codes. Digital signatures are used to authenticate the initial Diffie-Hellman key exchange. OTR supports two methods for authenticating each user's public signature key: Either by manually verifying the public key's fingerprint or by using the socialist millionaire protocol (SMP[25]), which verifies that both users share the same secret without revealing it. Finally, message authentication codes are used to authenticate chat messages and all following Diffie-Hellman key exchanges.

Deniability. OTR supports "weak" and "strong" (plausible) deniability. Message authentication codes ensure the "weak" deniability property because both parties share the same secret key, thus Alice could create fake messages in Bob's name. OTR uses a malleable encryption scheme (AES in counter mode) that allows anyone listing to the communication channel to forge messages. As a result, OTR also supports the "strong" deniable authentication property.

Anonymity. Anonymity is not a built-in security property of OTR. Higher level IM transport protocols such as XMPP can leak metadata information (e.g. current time and timezone, geolocation[26] and much more). In addition, OTR does not protect user identities against active attacks because long-term public keys, which clearly identify a person, can be revealed by a man-in-the-middle attack [23]. Moreover, OTR does not safeguard against timing and statistical attacks, which may allow an attacker to find out who is talking to whom and to guess the message content. However, using OTR on top of anonymity networks such as Tor could provide anonymity. For example, ChatSecure can access the Tor network by employing the proxy application Orbot[27].

Weaknesses. Even though OTR has strong security and privacy properties, several weaknesses exist. One disadvantage of the protocol is that it contains a reliability issue which is a design flaw. The reason for this is that OTR requires both parties to be online in order to perform the Diffie-Hellmann key exchange. As a result, messages might drop or are delivered unencrypted if one user goes offline. Another concern is that active attacks are possible if the authenticity of the client's public key is not correctly verified. For example, there is a plugin for the popular IM server software "ejabberd", which allows man-in-the-middle attacks if the public key's fingerprint is not accurately checked[28]. Apart from these weaknesses, the research community has discovered additional security flaws in the OTR protocol. Raimondo et. al. found several issues in OTR's authentication mechanism as well as replay vulnerabilities [13]. In addition, they revealed an UKS attack against OTR's key exchange procedure. Moreover, Bonneau et. al. discovered various attacks against OTR version 2.3 [6]. Firstly, they were able to perform a downgrade-attack in order to roll-back to a much weaker version of the protocol. Secondly, they showed attacks against the protocol's "strong" deniability property and integrity component. Finally, [26] presents weaknesses of OTR's "weak" deniability primitive.

3.2 TextSecure v2 Protocol

The open-source TextSecure v2 protocol[29] is based on OTR and developed by Open Whisper Systems[30]. The company also maintains a mobile Android/iOS application[31] that has been praised by Edward Snowden for its ease of use [16]. The TextSecure protocol has seen much popularity recently because it has been integrated into the famous IM application WhatsApp [49]. It has been also adapted by CyanogenMod which uses the protocol for its OS-level SMS provider in order to encrypt text messages [38]. The group-chat capable [39] protocol supports the following security properties:

Confidentiality. The end-to-end encryption mechanism is based on Curve25519 and AES-256 cryptographic

[19] https://jitsi.org/
[20] https://www.pidgin.im/
[21] https://adium.im/
[22] https://chatsecure.org/
[23] https://otr.cypherpunks.ca/Protocol-v3-4.0.0.html
[24] There is an OTR plugin for IRSSI: https://github.com/cryptodotis/irssi-otr.
[25] https://otr.cypherpunks.ca/help/authenticate.php
[26] http://xmpp.org/extensions/xep-0080.html
[27] https://guardianproject.info/apps/orbot/
[28] https://www.ejabberd.im/mod_otr
[29] https://github.com/WhisperSystems/TextSecure/wiki/ProtocolV2
[30] https://www.whispersystems.org
[31] There is also an unfinished browser version available at https://github.com/whispersystems/TextSecure-Browser.

primitives. TextSecure improves OTR's Diffie-Hellmann key exchange [35] by employing the Axolotl protocol[32].

Perfect Forward Secrecy. TextSecure supports perfect forward secrecy for asynchronous messaging [36] by generating a new session key for every new message.

Authentication. Long-term and ephemeral secret keys authenticate the triple Diffie-Hellman (DH) key exchange whereas short-term MAC (HMACSHA256) keys protect the actual chat messages.

Deniability The protocol uses an improved version of OTR's deniability property [37].

Weaknesses. The TextSecure protocol faces some weaknesses because it does not protect anonymity, thus allowing a third-party to collect metadata. Additionally, Bader et. al. showed an Unknown Key-Share Attack (UKS) against TextSecure [3]. Nevertheless, the researchers conclude that TextSecure's push messaging achieves the goals of authentication and confidentiality.

3.3 Silent Circle Instant Messaging Protocol

The Silent Circle Instant Messaging Protocol (SCIMP) is a secure open-source[33] IM protocol developed by Silent Circle[34]. The protocol is derived from ZRTP and enables private conversations over IM transports such as XMPP. Silent Circle has developed a mobile application called Silent Text 2.0 which is available for iOS and Android operating systems and has been audited independently [41]. SCIMP provides the following security features:

Confidentiality. SCIMP uses AES in counter with CBC-MAC (CCM) mode to provide strong encryption. The elliptic curve Diffie-Hellman key exchange is performed between the two endpoints only; thus end-to-end encryption is provided.

Perfect Forward Secrecy. Each message is encrypted with its own shared secret key, which is discarded after use.

Authentication. Before two users communicate, Silent Text 2.0 presents an "authentication string" which users can compare. If the initial communication was successful, all future sessions are authenticated via a cached "secret" [22]. Encrypting messages using AES in CCM mode ensures the secrecy as well as authenticity of chat messages.

Weaknesses. By design, SCIMP does not provide deniability (hence Bob can create fake messages in Alice's name). Another weakness is that SCIMP does not guarantee anonymous communication because metadata (leaked by IM transports) could be extracted in order to perform timeline analysis. Finally, a recent version of Silent Text contained a vulnerability, which allowed an attacker to decrypt messages [15].

3.4 Cryptocat

Cryptocat is an experimental open-source browser-based IM application which is available as a plugin for Chrome, Firefox, Safari and Opera. The JavaScript client is executed locally and communicates to a XMPP-BOSH[35] server via HTTPS. Cryptocat uses the OTR protocol for one-on-one encrypted communications and the mpOTR[36] protocol for multi-party encrypted group chat. Moreover, an iOS mobile application and a standalone Mac OS X application exist (an Android version is under current development). Cryptocat benefits from OTR's strong security and privacy properties:

- **Confidentiality.** It provides message confidentiality by using the OTR protocol for message encryption.
- **Perfect Forward Secrecy.** Session keys are discarded immediately after usage and are impossible to derive from any long-term key material.
- **Authentication.** In order to authenticate public signature keys, Cryptocat presents users with a fingerprint which can be manually verified [23].
- **Deniability.** Plausible deniability is guaranteed by the OTR protocol.

Weaknesses. On the downside, Cryptocat does not guarantee anonymity because it does not hide metadata[37] that is leaked by the XMPP protocol. The fact that Cryptocat is running in a browser is simultaneously its major strength and weakness. The advantage is that everyone has a browser and that there is no hurdle to install a separate application. However, a web browser is a popular attack vector and often contains software vulnerabilities; e.g. the famous pwn2own event regularly reveals exploitable browser bugs [8]. In addition, client-side JavaScript is not considered to be a secure programming language for implementing cryptography [1]. Furthermore, the results of commercial code audits have demonstrated several critical security flaws in Cryptocat's browser and iOS version [54] [14]; other vulnerabilities have been detected in older versions [53].

3.5 BitTorrent Bleep

Bleep[38] is a closed-source peer-to-peer application for internet telephony and asynchronous IM chat. The IM application is developed by BitTorrent Inc. and is available as an

[32] https://github.com/trevp/axolotl/wiki

[33] https://github.com/SilentCircle/silent-text

[34] Specification available at https://silentcircle.com/sites/default/themes/silentcircle/assets/downloads/SCIMP_paper.pdf.

[35] https://xmpp.org/about-xmpp/technology-overview/bosh/

[36] Goldberg et. al. proposed a multi-party variant of OTR called mpOTR [20].

[37] https://blog.crypto.cat/2013/06/cryptocat-who-has-your-metadata/

[38] http://labs.bittorrent.com/bleep/

alpha version for Mac OS X, Windows and Android. Bleep is based on the BitTorrent protocol, which uses a distributed hash table (DHT) to route messages. The IM application provides client (SIP UAC) and server (SIP server) functionalities. The SIP server builds the foundation of the decentralized peer-to-peer platform and serves as Bleep's back-end or engine. The SIP User Agent Client (UAC) provides the user interface and communicates to the SIP server. Bleep features the following security properties:

Confidentiality. Text messages are transmitted in an encrypted SIP tunnel over UDP between two peers. Bleep employs encryption protocols such as curve25519, ed25519 , salsa20, poly1305, and others.

Perfect Forward Secrecy. A new temporary session key is generated for each message and deleted immediately after usage [45].

Authentication. Bleep identifies users' by their public keys. The keys are stored on a directory server in order to allow lookups when a user wants to add a new contact. Bleep uses a third-party authority which verifies the authenticity of users' public keys [45].

Deniability. Bleep is still in a developing stage and lacks detailed documentation. Thus, we are uncertain whether Bleep offers deniability.

Anonymity. Bleep provides anonymity in the sense that there is no central repository for metadata storage and no central lookup service. In addition, Bleep makes use of the same DHT network as uTorrent, which already has millions of users. Another anonymity feature of Bleep is that it is difficult for a third-party to find out which public key corresponds to which user IP address [17].

Weaknesses. Even though Bleep uses the open-source cryptography library libsodium[39], the main source code remains confidential. Additionally, Bleep does not hide the endpoints' IP addresses, allowing an adversary to guess who is talking to whom. Moreover, Bleep routes clients through a relay server if a direct client-to-client communication is not possible. Hence, a compromised relay server could perform traffic analysis on the encrypted network packets. Finally, BitTorrent Inc. has been recently accused of adding "riskware" to their products; the company added a BitCoin mining application to uTorrent without explicit user consent[40].

4 Related work

Many other (secure) IM applications and protocols exist. For instance, there are proprietary IM applications which rely on a company-controlled client-server architecture: Viber, Telegram, Threema, Wickr and Surespot. Although Threema makes use of the open-source library NaCl, the protocol is kept confidential. Wickr features message destruction and a RSA-based key-exchange, thus providing perfect forward secrecy [22]. Nevertheless, its security design is weakly documented and there have not been any independent code reviews. Surespot[41] employs its own cryptographic protocol and is available as an Android and iOS application. However, there have not been any protocol audits either.

Anonymous peer-to-peer IM applications, which rely on a trustworthy distributed chat system, do not need any central servers. Several open-source solutions exist: For example, there is PyBitmessage[42], which uses the BitMessage[43] protocol in order to provide unobservability and untraceability. RetroShare, which is a software bundle for private communications, also supports private IM in a friend-to-friend network scenario. Then there is Ricochet[44], which is a peer-to-peer IM system based on Tor hidden services. Other decentralized IM applications include Tox[45], Pond[46], or psyced[47].

Finally, there are secure IM protocols which have not been widely adopted: SILC[48], IMKE [34] and FiSH[49]. Nevertheless, there are SILC plugins for Pidgin and Irssi; FiSH enables encrypted IRC chat via Blowfish with pre-shared keys.

5 Discussion

This section compares the secure IM technologies that have been presented in this survey paper. For the comparison, two different metrics are applied: One for secure IM protocols and one for secure IM applications. The metrics of the protocol comparison include [19]:

- **End-to-end encryption.** Whether the key-exchange occurs between the two endpoints only.
- **PFS.** Whether the protocol supports Perfect Forward Secrecy.
- **Authentication.** Whether participants can be certain about the identity of correspondents.
- **Deniability.** Whether Bob is not able to proof to a third-party that Alice sent a certain message.
- **Recent review.** Whether there has been a recent independent protocol review.
- **Documentation.** Whether the cryptography design is well documented.

[39] https://github.com/jedisct1/libsodium

[40] http://forum.utorrent.com/topic/95041-warning-epicscale-riskware-installed-with-latest-utorrent/

[41] https://www.surespot.me/

[42] https://github.com/Bitmessage/PyBitmessage

[43] https://bitmessage.org/bitmessage.pdf

[44] https://github.com/ricochet-im/ricochet

[45] https://tox.im/

[46] https://pond.imperialviolet.org/

[47] http://www.psyced.org/

[48] http://tools.ietf.org/id/draft-riikonen-silc-spec-09.txt

[49] http://ultrx.net/doc/fish/

[50] Diffie-Hellman key exchange.

[51] e.g. https://www.ejabberd.im/

[52] https://github.com/WhisperSystems/TextSecure-Server

[53] https://github.com/cryptocat/cryptocat/wiki/Server-Deployment-Instructions

	End-to-end encryption	PFS	Authentication	Deniability	Recent review	Documentation
TLS	no	DH[50]	X.509	MACs	no	+
PGP	3DES/RSA	no	WOT	no	no	+
OTR	AES-CTR	DH	Fingerprint/SMP	MACs/AES-CTR	yes	++
TextSecure	AES-CTR	ECDH	Fingerprint	MACs/AES-CTR	yes	++
SCIMP	AES-CCM	ECDH	SAS	no	yes	++

Table 1: Evaluation matrix of IM protocols.

	Protocol(s)	Metadata protection	Open-Source	Recent Code Audit	Documentation	Stable software version
Pidgin	OTR	−−	client & server (XMPP[51])	yes	++	yes
TextSecure	TextSecure	−−	client & server[52]	yes	++	yes
Silent Text	SCIMP	−−	client only	yes	+	yes
Cryptocat	OTR	−−	client & server (XMPP[53])	yes	++	no (experimental)
Bleep	BitTorrent/SIP/RTP	+	no	no	−−	no (alpha)

Table 2: Evaluation matrix of secure IM applications.

Table 1 summarizes the comparison of secure IM protocols. The major findings are that TLS and PGP should not be used for secure instant messaging because they lack several security features. The OTR protocol guarantees strong "face-to-face privacy", but does not provide asynchronous messaging and anonymity. Nevertheless, most of the OTR capable applications support location anonymity by routing IM traffic through the Tor network. The TextSecure protocol includes OTR's features, but does not protect the user's metadata. SCIMP has similar strong security properties, but lacks deniability and anonymity features.

In order to compare secure IM application, slightly different metrics apply:

- **Protocol(s).** Which protocol(s) does the application rely on.
- **Metadata protection.** Whether the application protects the user's metadata.
- **Open-Source.** Whether the client and server application is open for public review.
- **Recent code audit.** Whether there has been a recent independent security audit.
- **Documentation.** Whether the application's technical components are well documented.
- **Stable software version.** Whether the current software version is stable.

Table 2 illustrates a comparison of the IM applications that have been discussed. One of the major results of the analysis is that TextSecure and Silent Text enable secure instant messaging. Nevertheless, both applications do not hide the user's metadata and cannot be used via the Tor network. Cryptocat is a secure IM application that is easy to use and heavily audited by third-parties. However, Cryptocat is an experimental application and is executed by insecure web browsers. Bleep is the most anonymous IM solution presented in this survey. However, it is closed-source and currently an alpha version.

In summary, we recommend the use of TextSecure for secure instant messaging as it features all aspects of a private face-to-face conversation. Additionally, TextSecure works well on mobile devices, thus providing asynchronous IM.

6 Conclusion

This seminar paper explored few popular secure IM protocols and applications; to mention all available secure instant messaging services is beyond the scope of this work. This is actually a problem for the normal user, as he or she does not know which application to choose. For this reason, most people use the mainstream less secure IM services and also because most of their contacts already use the same IM platform. Another major challenge for secure IM applications is the trade-off between usability and security. Most of the insecure IM applications are easy to use because they do not rely on security best practices. On the other hand, secure instant messaging applications can be hard to use for the everyday user, e.g. verifying OTR fingerprints can be difficult. Another issue is habituation, i.e. when users accept security warnings without reading them carefully.

Finally, no security system is fullproof. One cannot prevent targeted attacks such as backdoors in a user's hardware or software; for example there is a huge market for Android/iOS exploits with market prices above 500.000 $. To

quote Matthew Green [22]:

> The real issue is that they [secure IM applications] each run on a vulnerable, networked platform.

References

[1] Javascript cryptography considered harmful. http://matasano.com/articles/javascript-cryptography/, April 2013. Accessed: 2015-03-21.

[2] R. Andrews. Exploring encrypted skype conversations, in clear-text. https://www.bluecoat.com/security-blog/2014-01-02/exploring-encrypted-skype-conversations-clear-text, January 2014. Accessed: 2015-02-24.

[3] T. F. C. M. C. Bader and F. B. J. S. T. Holz. How secure is textsecure?

[4] M. Barakat. Jeffrey sterling, ex-cia officer, convicted of leaking secrets to reporter. http://www.washingtontimes.com/news/2015/jan/26/deliberation-to-reach-third-day-in-cia-leak-case/, January 2015. Accessed: 2015-03-11.

[5] D. Bean. How to save snapchat pictures without the sender knowing (shhhh). https://www.yahoo.com/tech/how-to-save-snapchat-pictures-without-the-sender-80875346215.html, March 2014. Accessed: 2015-03-20.

[6] J. Bonneau and A. Morrison. Finite-state security analysis of otr version 2.

[7] N. Borisov, I. Goldberg, and E. Brewer. Off-the-record communication, or, why not to use pgp. In *Proceedings of the 2004 ACM workshop on Privacy in the electronic society*, pages 77–84. ACM, 2004.

[8] C. Brook. All major browsers fall at pwn2own day 2. https://threatpost.com/all-major-browsers-fall-at-pwn2own-day-2/111731, March 2015. Accessed: 2015-03-21.

[9] I. Brown, A. Back, and B. Laurie. Forward secrecy extensions for openpgp. Internet-draft, The Internet Engineering Task Force, October 2001. http://tools.ietf.org/html/draft-brown-pgp-pfs-03.

[10] C. Cattiaux. imessage privacy. http://blog.quarkslab.com/imessage-privacy.html, October 2013. Accessed: 2015-03-20.

[11] S. E. Coull and K. P. Dyer. Traffic analysis of encrypted messaging services: Apple imessage and beyond. *ACM SIGCOMM Computer Communication Review*, 44(5):5–11, 2014.

[12] M. Di Raimondo and R. Gennaro. New approaches for deniable authentication. In *Proceedings of the 12th ACM conference on Computer and communications security*, pages 112–121. ACM, 2005.

[13] M. Di Raimondo, R. Gennaro, and H. Krawczyk. Secure off-the-record messaging. In *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, pages 81–89. ACM, 2005.

[14] A. Diquet, D. Thiel, and S. Stender. Open technology fund cryptocat ios. Technical report, iSEC Partners, February 2014. http://isecpartners.github.io/publications/iSEC_Cryptocat_iOS.pdf.

[15] M. Dowd. Blackpwn: Blackphone silenttext type confusion vulnerability. http://blog.azimuthsecurity.com/2015/01/blackpwn-blackphone-silenttext-type.html, January 2015. Accessed: 2015-03-13.

[16] M. Eddy. Snowden to sxsw: Here's how to keep the nsa out of your stuff. http://securitywatch.pcmag.com/security/321511-snowden-to-sxsw-here-s-how-to-keep-the-nsa-out-of-your-stuff, March 2014. Accessed: 2015-03-19.

[17] F. Fadaie. How does bleep work? http://engineering.bittorrent.com/2014/09/17/how-does-bleep-work/, September 2014. Accessed: 2015-03-21.

[18] E. F. Foundation. 20131014-wapo-content acquisition optimization2. https://www.eff.org/document/2013-10-14-wapo-content-acquisition-optimization2, 2013. Accessed: 2015-02-03.

[19] E. F. Foundation. Secure messaging scorecard. https://www.eff.org/secure-messaging-scorecard, March 2015. Accessed: 2015-03-30.

[20] I. Goldberg, B. Ustaoğlu, M. D. Van Gundy, and H. Chen. Multi-party off-the-record messaging. In *Proceedings of the 16th ACM conference on Computer and communications security*, pages 358–368. ACM, 2009.

[21] M. Green. Dear apple: Please set imessage free. http://blog.cryptographyengineering.com/2012/08/dear-apple-please-set-imessage-free.html, August 2012. Accessed: 2015-03-20.

[22] M. Green. Here come the encryption apps! http://blog.cryptographyengineering.com/2013/03/here-come-encryption-apps.html, March 2013. Accessed: 2015-03-20.

[23] M. Green. Noodling about im protocols. http://blog.cryptographyengineering.com/2014/07/noodling-about-im-protocols.html, July 2014. Accessed: 2015-03-10.

[24] M. Green. What's the matter with pgp? http://blog.cryptographyengineering.com/2014/08/whats-matter-with-pgp.html, August 2014. Accessed: 2015-03-13.

[25] A. Greenberg. Whatsapp just switched on end-to-end encryption for hundreds of millions of users. http://www.wired.com/2014/11/whatsapp-encrypted-messaging/, 2014. Accessed: 2015-02-03.

[26] greg. Secure function evaluation vs. deniability in otr and similar protocols. http://phrack.org/issues/68/14.html#article, 2012.

[27] C. G. Günther. An identity-based key-exchange protocol. In *Advances in CryptologyâĂŤEurocryptâĂŹ89*, pages 29–37. Springer, 1990.

[28] N. Hindocha and E. Chien. Malicious threats and vulnerabilities in instant messaging. In *Virus Bulletin Conference, vb2003*, 2003.

[29] G. Inc. Transparency report. https://www.google.com/transparencyreport/userdatarequests/, 2014. Accessed: 2015-03-03.

[30] M. Jakobsson and M. Yung. Proving without knowing: On oblivious, agnostic and blindfolded provers. In *Advances in CryptologyâĂŤCRYPTOâĂŹ96*, pages 186–200. Springer, 1996.

[31] lynX. 15 reasons not to start using pgp. http://www.secushare.org/PGP, July 2014. Accessed: 2015-03-13.

[32] M. Mannan and P. C. van Oorschot. Secure public instant messaging: A survey. *Proceedings of Privacy, Security and Trust*, 2004.

[33] M. Mannan and P. C. van Oorschot. On instant messaging worms, analysis and countermeasures. In *Proceedings of the 2005 ACM workshop on Rapid malcode*, pages 2–11. ACM, 2005.

[34] M. Mannan and P. C. Van Oorschot. A protocol for secure public instant messaging. In *Financial Cryptography and Data Security*, pages 20–35. Springer, 2006.

[35] M. Marlinspike. Advanced cryptographic ratcheting. https://www.whispersystems.org/blog/advanced-ratcheting/, November 2013. Accessed: 2015-03-19.

[36] M. Marlinspike. Forward secrecy for asynchronous messages. https://whispersystems.org/blog/asynchronous-security/, August 2013. Accessed: 2015-03-19.

[37] M. Marlinspike. Simplifying otr deniability. https://whispersystems.org/blog/simplifying-otr-deniability/, July 2013. Accessed: 2015-03-19.

[38] M. Marlinspike. Textsecure, now with 10 million more users. https://whispersystems.org/blog/cyanogenintegration/, December 2013. Accessed: 2015-03-19.

[39] M. Marlinspike. Private group messaging. https://whispersystems.org/blog/private-groups/, May 2014. Accessed: 2015-03-19.

[40] M. Marlinspike. Gpg and me. http://www.thoughtcrime.org/blog/gpg-and-me/, February 2015. Accessed: 2015-03-13.

[41] J. Matonis. Pressure increases on silent circle to release application source code. http://www.forbes.com/sites/jonmatonis/2013/02/06/pressure-increases-on-silent-circle-to-release-application-source-code/, June 2013. Accessed: 2015-03-20.

[42] T. McCarthy. Nsa director defends plan to maintain 'backdoors' into technology companies. http://www.theguardian.com/us-news/2015/feb/23/nsa-director-defends-backdoors-into-technology-companies, February 2015. Accessed: 2015-02-24.

[43] G. News and M. Limited. A guardian guide to your metadata. http://www.theguardian.com/technology/interactive/2013/jun/12/what-is-metadata-nsa-surveillance, June 2013. Accessed: 2015-03-11.

[44] G. News and M. Limited. Project bullrun – classification guide to the nsa's decryption program. http://www.theguardian.com/world/interactive/2013/sep/05/nsa-project-bullrun-classification-guide, 2013. Accessed: 2015-02-03.

[45] A. Norberg. Authentication and forward secrecy in bleep. http://engineering.bittorrent.com/2014/12/11/authentication-and-forward-secrecy-in-bleep/, December 2014. Accessed: 2015-03-21.

[46] J. S. Park and T. Sierra. Security analyses for enterprise instant messaging (eim) systems. *Information Systems Security*, 14(1):26–39, 2005.

[47] J. Scahill and G. Greenwald. The nsaâĂŹs secret role in the u.s. assassination program. https://firstlook.org/theintercept/2014/02/10/the-nsas-secret-role/, February 2014. Accessed: 2015-03-11.

[48] B. Schneier. New details on skype eavesdropping. https://www.schneier.com/blog/archives/2013/06/new_details_on.html, 2013. Accessed: 2015-02-03.

[49] B. Schneier. Whatsapp is now end-to-end encrypted. https://www.schneier.com/blog/archives/2014/11/whatsapp_is_now.html, November 2014. Accessed: 2015-03-19.

[50] S. Schoen. Why doesn't skype include stronger protections against eavesdropping? https://www.eff.org/deeplinks/2013/07/why-doesnt-skype-include-stronger-protections-against-eavesdropping, July 2013. Accessed: 2015-02-24.

[51] K. Shubber. Blackberry gives indian government ability to intercept messages. http://www.wired.co.uk/news/archive/2013-07/11/blackberry-india, 2013. Accessed: 2015-03-03.

[52] Systems and N. A. Center. Secure instant messaging. Technical report, National Security Agency, 2007.

[53] S. Thomas. Decryptocat. http://tobtu.com/decryptocat.php. Accessed: 2015-03-21.

[54] N. Wilcox, Z. Wilcox, D. Hopwood, and D. Bacon. Report of security audit of cryptocat. Technical report, Least Authority, April 2014. https://leastauthority.com/static/publications/LeastAuthority-Cryptocat-audit-report.pdf.

A survey on performance of SfM and RGB-D based 3D indoor mapping

Junyang Shi
Student number: 466958
junyang.shi@aalto.fi

Abstract

Extracting an indoor map from 3D models is a new trend for building maps for indoor environments. Modern computer vision techniques makes it possible to reconstruct a 3D model from images. This paper surveys two main computer vision techniques regarding 3D indoor mapping call SfM and RGB-D. It analyse the design and performance of SfM and RGB-D mapping systems based on current literature. Then, a comparison between SfM and RGB-D is made focuses on properties of applicable environment, system input, accuracy and costs. It concludes that RGB-D mapping systems have relatively high accuracy and high costs in indoor environment compared to SfM based mapping systems. Current work limitation and future challenges are presented.

KEYWORDS: SfM, RGB-D, 3D indoor mapping, image collections

1 Introduction

Google Maps and GNSS have revolutionized the way we live and travel. Conventionally, these maps are built by surveying the world around us, but indoor environments such as offices, museums, and airports can be very time consuming to survey due to occlusions. Moreover, due to time constraints and obstruction of satellite signals, maps for indoor environments are difficult to build.

The development of three-dimensional modeling of building interiors make it possible for indoor maps to be extracted from 3D models. The 3D models of objects, structures and buildings are introduced for use in a variety of applications, such as virtual museums, counteracting terrorism, urban planning, simulation for disaster management, navigation systems, the mobile service industry [8] and many other applications.

Traditionally, 3D reconstruction is conducted by photogrammetry [17] which results in huge time and financial costs. Computer vision technologies, such as Structure-from-Motion (SFM), RGB-D mapping [7] and Simultaneous Localization and Mapping [1] have been applied to reconstruct 3D modeling systems in recent years. All of the computer visual techniques metioned above reconstruct 3D models without demanding photogrammetry information, thus achieving comparative efficiency in computation.

Recovering 3D structures from a set of 2D images has become very popular and several research groups have studied the field over the past few years [11, 14, 13]. Photo Tourism [11] typically shows how 3D structures can be recovered through photo matching and analyzing. The research work [14] explains the possibility to reconstruct 3D models from Internet photos. Moreover, the research work [13] shows promising results for image-based 3D reconstruction on city-scale outdoor scenes.

Researchers usually call 3D models that are built by computer vision techniques a point cloud. A point cloud is a set of 3D points which represent the shape of a 3D object. Each 3D point has 3(x,y,z) coordinates, corresponding color information of the point as well as a list of measurements used to estimate the position of the point. Both sparse and dense point clouds exist.

This paper makes a survey on performance of two main computer vision techniques regarding 3D mapping from images named SfM and RGB-D. We firstly analyse some interesting 3D mapping systems and their experiment performance based on SfM and RGB-D. Then we make a comparison between SFM and RGB-D techniques with focus on applicable environment, system input, accuracy and costs. Finally, we illustrate future challenges in image-based 3D mapping field.

The remainder of this paper is structured as follows. Section two presents an overview of the Structure-from-Motion and RGB-D mapping techniques. Section 3 surveys current research works. Section 4 makes a comparison between SfM and RGB-D techniques based on research works. Current work limitation and future challenges are illustrated in Section 5. Finally, section 6 concludes the paper.

2 Background

2.1 Structure from Motion

Structure from Motion (SfM) is a major approach to build 3D models from images. SfM was introduced to search for the same feature points in photo collections, which recovers the 3D position that correspond to objects seen in photos as well as estimate the location and orientation of a camera [13]. After determining the camera orientation parameters, the 3D coordinates of the cameras and the image-based point cloud can be determined. A number of tools are created for generating point cloud from 2D images both by commercial vendors and research groups.
Typically, SfM methods comprises several steps, which consists of feature extraction, feature correspondence, skeletal

set selection and bundle adjustment [16]. Feature extraction solves the problem of detecting distinctive, repeatable features for matching in an image. Feature correspondence verifies(via feature matching) across input images in order to estimate a single 3D point from all the features. Bundle Adjustment finds the correct positions of the 3D points and recovers camera parameters. Skeletal set selection is a supplementary procedure which identifies a subset of images to be reconstructed in order to improve reconstruction performance.

2.1.1 Feature Extraction

One of the most successful feature extraction algorithms is Scale Invariant Feature Transform (SIFT) [9]. SIFT features are invariant and it works well with changes in scale and rotation across input images. However, the performance of feature matching based on the SIFT algorithm can be enhanced since SIFT features do not contain the location of other features in the image. With respect to improving speed and accuracy, several enhancements including PCA-SIFT, CSIFT, ASIFT, SURF and other feature extraction algorithms have been proposed [12].

2.1.2 Feature Correspondence

Feature correspondence starts with matching features across image pairs. Typically, the features of one image are stored in the form of a tree abstract data structure which provides nearest neighbor search [11]. And the features from the other image are used to query stored features and obtain the nearest neighbor which is then declared a match. After obtaining pair matching information, the matched features are grouped to build a features database [12]. Subsequent features are matched by querying feature database.

2.1.3 Bundle Adjustment

Recovering the correct position of the 3D points as well as camera parameters requires solving a non-linear optimization problem. Generally, bundle adjustment is a process which minimizes the total squared reprojection error across multiple overlap image. Since directly solving the optimization problem which is a collinearity formula is hard, most sfm systems implement an incremental approach on unordered photo collections, i.e. starting with a small reconstruction, then growing a few images at a time. The process is repeated until no more images remain [13]. However, the incremental approach indeed consumes a large amount of time and computation resources. The following section shows a good solution to large scale photo collections.

2.1.4 Skeletal Set Selection

Photo collections, especially those found on Internet photo-share sites, by their nature are unordered and highly redundant. A lot of such photos are taken from nearby viewpoints and processing all of them does not necessarily add to the reconstruction [13]. Moreover, sfm scaling techniques tend not to apply well to irregular photo collections. Therefore, it is necessary to select a smaller set of images which capture essential geometry information of the scene. Researchers called the selected subset of views a skeletal set. Once the skeletal set is identified, the reconstruction process follows into two steps. First, conducting the reconstruction process on the skeletal images. Then, the remaining photos will be added to the reconstruction by estimating each camera's pose with respect to known 3D points matched to that image [13] based on scalable algorithms. As shown in [15], this process results in an order of magnitude and increases the computaion efficiency with little or no loss of accuracy.

2.1.5 Indoor Map Extraction

Indoor map needs to be extracted after obtaining 3D point cloud built from SfM. The main challenge of recovering indoor map is to extract the shape of indoor obstacles out of a 3D point cloud.

One straightforward way proposed in [2] contains two steps. Firstly, identifying points which represent obstacles in the 3D point. In detail, a flatten process is conducted in which a threshold is set to filter floor points and other points that represent obstacles are assigned a zero depth coordinate thus being able to be placed on a ground plane. Subsequently, the solution proposed by Duckham [4] which form a shape from points is used to generate non-convex polygons from unordered 2D points in the ground plane. Finally, an indoor map is built through extruding the shape of obstacles towards the depth coordinate and placing them on a ground plane [2].

Another approach is to update incomplete 3D model produced by Multi-view stereo(MVS) reconstruction algorithm. MVS is one of the most successful dense approach which recovers dense and accurate 3D model by processing the sparse point cloud built through bundle adjustment. The key idea is to compute a full model through some form of interpolation after MVS reconstruction. Manhattan-world Stereo [5] is a popular algorithm to compute a full model. It exploits the Manhattan-world assumption – surfaces are piece-wise planar and aligned with three dominant directions [5]. This assumption is to solve the challenges of indoor environments including low lighting and textureless surfaces. At the end, a complete depth map for every input image will be generated. The depth maps are seen as input to be merged into a 3D mesh, i.e. an indoor map.

2.2 RGB-D Mapping

RGB-D 3D mapping uses color frames produced by RGB-D cameras to generate dense 3D models of building interiors. RGB-D cameras provides RGB color information along with per-pixel depth information of images. Henry et al [6] showed that such cameras are suitable for dense 3D modeling. Specifically, RGB-D cameras are able to capture mid-resolution depth and appearance information at high data rates, typically 30 frames per second [7].

In RGB-D 3D mapping process, the main steps are similar to SfM. First of all, feature points contained in the RGB images are extracted. Then, featrues are matched through a random sample consensus (RANSAC) procedure [7]. Subsequently, the step is named loop closure detecting [16] which iteratively matches the current color frames to a subset of

registed frames that resulting from feature matches. It is used to determine the best alignment between frames and to check if the current frame is a revisit to a known scene [3] as well. Finally, RGB-D mapping builds a global model using small planar colored surface patches [7].

The main technical difference for RGB-D 3D modeling compared to SfM is that RGB-D cameras are able to capture rich visual information of a particular scene, while 3D point clouds built from SfM fails to utilize valuable information contained in images. Nevertheless, the robustness for use of depth cameras has great potential to be enhanced. The RGB-D cameras are developed by a number of commercial vendors such as Microsoft and PrimeSense.

3 Survey of Related Works

Reconstructing building interiors based on computer vision techniques is a huge challenge. The indoor environment is always interconnected resulting in only a subset of entire objects being visable in each image, thus impacting the feature matching procedure in SfM. Moreover, building interiors are full of painted walls and other textureless objects which lack of distinctive key features. Because of the fact that MVS algorithm performs relatively poorly for texture-poor objects, more efficient approaches are needed with repect to recovering building interiors from images.

3.1 SfM

3.1.1 Annotated Map Approach

Ricardo el al [10] aim to map large indoor spaces using Internet photos. Since image collections downloaded from Internet fail to cover the global view of a particular scene, the 3D models generated through SfM are disjointed and incomplete. Ricardo el al [10] present an approach to recover the global layout of the 3D model reconstructed out of Internet photos with the help of a provided map. The key problem that exists in their approach is to extract and integrate the position, orientation and scale cues from 3D models and annotated map of sites. Generally, their process comprises two steps as follows.

Firstly, 3D models are assigned to the 2D regions of sites in a global frame, i.e. locating the position of 3D models. The 2D regions are recovered from an annotated map provided. In detail, the annotated map of a site is firstly found online, then a semi-automatic method is implemented which extracts the spatial layout of the objects on the map. As a result, the corresponding 2D region collections of a given site are recovered, such as rooms, and hallways from the floorplans. Subsequently, a set of discrete 3D models reconstructed by SfM are placed into the 2D region using cues from Google Image Search and the shape of the rooms. As a result, an indoor map the global layout of the 3D models is completed.

Secondly, the orientation of 3D models needs to be determined after being placed into the 2D region. The global layout of the 3D models fails to consider the orientation of the 3D models, thus yielding an incomplete indoor map. Ricardo el al [10] introduce a novel crowd flow cue to measure how people move across the site to recover 3D geometry orientation. Based on the proposed crowd flow cue, travel paths of people are measured and then indicate concrete information on the orientation of 3D models. In addition, Ricardo el al [jigsaw] create interactive navigation with respect to using the map. In the navigation mode of the map, when a user clicks on a room that a 3D model has been placed in, the viewpoint of user is directly directed into the aligned 3D model of the room. Fig. 1 shows promising results of their experiments [10] on major tourist sites.

3.1.2 Manhanttan-world Stereo Approach

Researchers in [5] proposed an automated computer vision system to reconstruct an entire floor of a house. The presented final 3D models [5] are impressive (see Fig. 2). The proposed system is fully automated and utilizes a number of exsiting techniques including SfM, MVS algorithm and manhattan-world stereo algorithm. Basically, the system consists of four steps. Firstly, SfM technique and MVS dense approach are used to compute the location, orientation and parameters of cameras for input image collections. Secondly, the system implements the Manhattan-world Stereo algorithm to generate depth maps from oriented 3D points after MVS reconstruction. Subsequently, based on a depth map integration approach, the input depth maps produced a axis aligned 3D mesh model. Finally, users are allowed to explore the reconstructed 3D environment with an image-based, interactive 3D viewer.

SfM, MVS algorithm and Manhattan-world Stereo algorithm have already been discussed in Section 2. The depth map integration approach used aims to produce a single 3D mesh model out of depth maps generated by Manhattan-world Stereo algorithm. Due to the scalability challenge of building interior objects, it is not practical to reconstruct detailed and complex 3D models. Therefore, the depth map integration approach focuses on extracting a simplified, axis-aligned 3D mesh model. Specifically, the integration algorithm merges the depth maps into a surface based on a set of complex formulations [5].

The 3D viewer aims to visualize the reconstructed indoor environment. It takes image collections and the axis-aligned 3D mesh model as input. It including two modes and is able to keep track of distances to all the cameras. Users can navigate either the input images or the rebuilt 3D space. Specifically, the viewpoint of the user is always set to the closest camera combined with data such as current viewing positions and directions.

3.2 RGB-D mapping

3.2.1 Interactive System Approach

Hao Du et al [3] proposed a prototype mobile system which builds dense and complete 3D models out of images for indoor environment. The system relies on depth cameras from which color and depth frames are produced.

Generally, the system is designed to run on a laptop and take in real-time images from depth cameras hold by user. It is an interactive approach that utilizes human input to cope

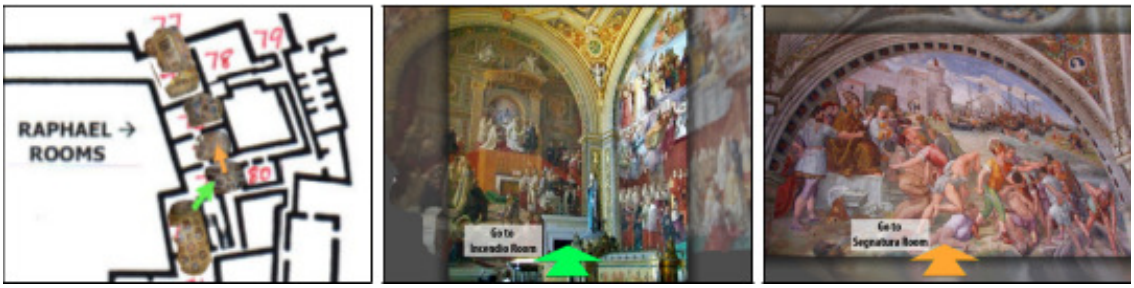

Figure 1: The 3D jigsaw puzzle: interactive visualization [picture from [10]]

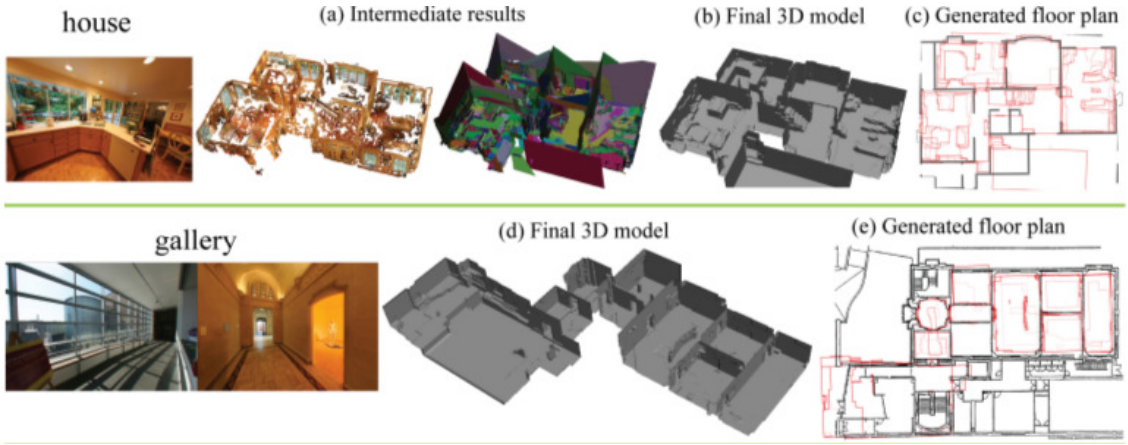

Figure 2: The final 3D models of house and gallery [picture from [5]]

with textureless and low lighting indoor environment. It offers real-time visualization of the reconstructed indoor environment to user, i.e. user can view how the 3D map grows in real-time on laptop. Moreover, the system provides feedback for users to change viewpoint of depth cameras in order to achieve robustness and completeness of the 3D map. Additionally, due to the utilization of depth sensor in depth cameras, the reconstructed 3D model are competitively dense compared to other existing 3D modeling system. The prototype mobile system is easy to use for both expert or non-expert.

The most advanced feature of the system is that it offers user real-time interaction. In detail, the system automatically provides suggestions on where the map may be incomplete. Therefore, with real-time guidance, user can align speed of motion and viewpoints of cameras in order to cover incomplete areas. Moreover, the system guarantees that frames which do not match recent frames will not be registered into the map, thus achieving registrating consecutive frames. In addition, the system is able to detect whether user is revisiting a known scene through matching "keyframes" and user is allowed to control corresponding detect algorithm. This function is named "loop closure" which further guarantee the system to yield globally consistent maps. In a word, the robustness and completeness of the prototype system is achieved to a high extent.

The promising result of the prototype mobile system [3] is showed fig. 3. Such a detailed indoor 3D map can be used for localization, navigation, measuring and other possible directions.

4 Comparison

This section conducts a comparison between SfM and RGB-D mapping techniques from different perspective named applicable environment, 3D mapping system input, accuracy and costs based on above research works. The table 1 presents the general comparison result.

4.1 Applicable environment

SfM is widely used for both indoor and outdoor spaces. However, SfM works relatively poor in indoor environment due to textureless objects such as walls which results in limited number of distinctive features. The main principle of the SfM is to establish the relationship between the different images [17]. Fortunately, SfM is easily to be combined with other computer vision techniques such as MVS or Manhatta-world stereo algorithms so as to improve 3D modeling accuracy.

RGB-D is mainly targeted at building interior. It is not practical for RGB-D to focus on outdoor space since it requires global consecutive frames from a particular scene. Usually,

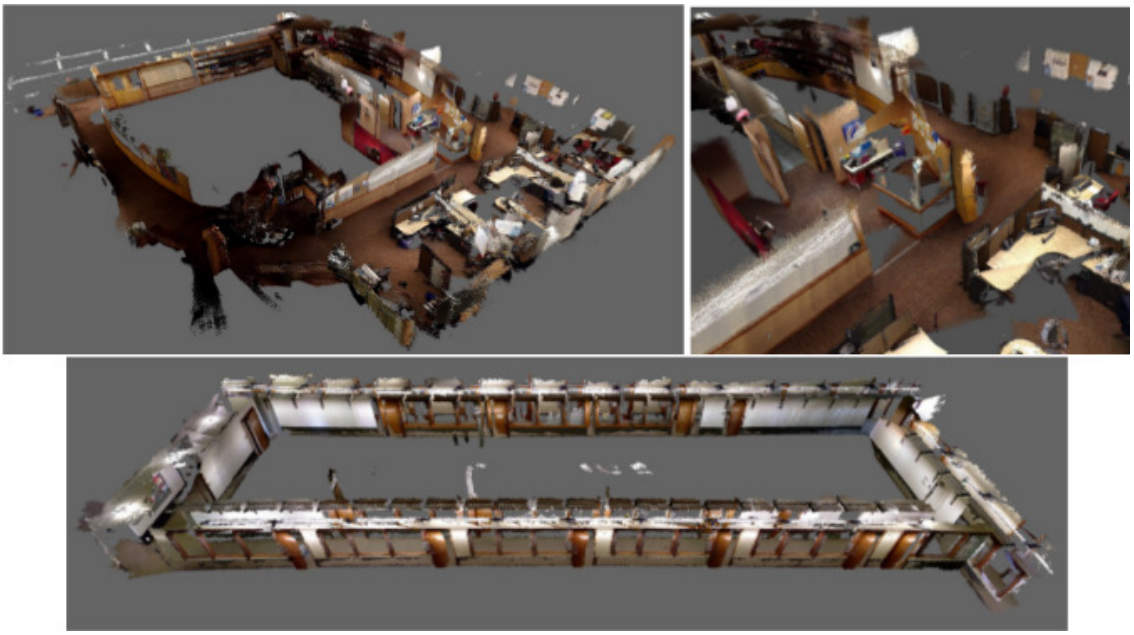

Figure 3: A variety of indoor spaces captured [picture from [3]]

RGB-D mapping systems are interactive and takes in real-time photographs from depth camera. Further more, RGB-D mapping system along with SLAM mapping system can be applied to robotics which estimates the robot pose and the scene geometry.

4.2 3D mapping system input and algorithm

SfM takes normal photo collections as mapping system input. Those images can be obtained from a moving camera or Internet photo sharing sites. Apart from basic feature extraction, feature matching and bundle adjustment algorithms, other algorithms includes MVS algorithms, stereo algorithm and depth-map integration algorithm are used to generate dense 3D indoor map.
RGB-D usually uses consecutive frames taken directly from depth cameras in a particular building interior. Usually RGB-D mapping system has more complicated algorithms than SfM since RGB-D deals with color and depth information. Corresponding alignment algorithms and combinations such as 3-point matching algorithm implemented in [3] uses both depth and color information [6]. Additionally, automatic matching algorithms needs to be implemented to provide frame registration filtering and interactive loop closure function.

4.3 Accuracy

3D mapping system based on SfM along tends to achieve sparse 3D point cloud in indoor environment. And when using images from Internet photo sharing sites, lack of consecutive images results in low completeness indoor map. The input images also have an impact on accuracy since nosiy image search results or lack of image cues render misplacing 3D models [10]. Nevertheless, combining with Manhatta-world stereo algorithm presented in [5], SfM mapping system shows promising performance in handle large-scare indoor scene. However, Manhatta-world stero algorithm consumed large runtime, therefore challenges still remains to improve the speed of SfM mapping system. RGB-D mapping system achieves relatively high accuracy. The indoor maps presented in [3] shows rich details for both large-size or small-size indoor space. The advantage of RGB-D mapping system is that the depth information provided by consumer depth cameras reduces reprojection errors, thus achieving promising results.

4.4 Costs

Generally, the more large scare of indoor scene targeted, the more computation resources are required, the higher costs are generated. For the same scale of building interior, RGB-D mapping system has relatively higher costs. Even if a low-cost RGB-D camera such as Microsoft Kinect is used, RGB-D mapping system requires consecutive frames taken from depth camera while SfM only demands unordered photo collections originated from Internet photo sharing sites. In addition, the costs also depends on the complexity of the 3D mapping system. Some 3D mapping systems implement a number of algorithms in order to achieve high accuracy. In this case, such as presented in [5], large amount of runtime are consumed and high costs are produced.

5 Challenges

Despite the promising results presented by some research groups regarding image-based 3D mapping, several chal-

Properties/Techniques	SfM	RGB-D
Applicable Environment	indoor and outdoor	mainly indoor
System Input	unordered photo collections	consecutive depth frames
Accuracy	relatively low	relatively high
Costs	relatively low	relatively high

Table 1: Comparison of SfM and RGB-D

lenges remain to be solved as follows.

- When using SfM with Internet photos as system input, missing focal length information of cameras may result in biased results in the process of mapping.
- Model updating is needed since indoor scenes such as the layout of a coffee room or interior decoration of a shopping center may change from time to time.
- Advanced computing approaches that reduce time consuming need to be developed for larger-scale scenes such as whole building interiors that comprise multiple floors.
- Accurate and detailed methods are needed to handle irregular indoor structures such as non-axis aligned surfaces.
- Exploring applications in depth is necessary to effectively use the detailed indoor maps.

6 Conclusion

This paper surveys two main computer vision techniques SfM and RGB-D which reconstruct 3D indoor map from images. In this paper, we have analysed the design and performance of 3D mapping systems using SfM and RGB-D based on current research works. Our comparison results have shown RGB-B mapping systems have relatively strict input requirement while SfM only demands unordered image collections. SfM are more widely used in both indoor and outdoor spaces while RGB-D is mainly used in indoor environment. Generally, RGB-D mapping system has relatively high accuracy as well as high cost. SfM performances relatively poor in textureless indoor environment. Nevertheless, mapping systems use SfM are able to improve accuracy via implementing further algorithms. And challenges includes reducing runtime and costs remain for both SfM and RGB-D 3D mapping systems.

Reference

[1] A. J. Davison, I. D. Reid, N. D. Molton, and O. Stasse. Monoslam: Real-time single camera slam. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 29(6):1052–1067, 2007.

[2] J. Dong, Y. Xiao, and A. Ylä-Jääski. Exploring the potential of indoor mapping and localization using internet photos.

[3] H. Du, P. Henry, X. Ren, M. Cheng, D. B. Goldman, S. M. Seitz, and D. Fox. Interactive 3d modeling of indoor environments with a consumer depth camera. In *Proceedings of the 13th international conference on Ubiquitous computing*, pages 75–84. ACM, 2011.

[4] M. Duckham, L. Kulik, M. Worboys, and A. Galton. Efficient generation of simple polygons for characterizing the shape of a set of points in the plane. *Pattern Recognition*, 41(10):3224–3236, 2008.

[5] Y. Furukawa, B. Curless, S. M. Seitz, and R. Szeliski. Reconstructing building interiors from images. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 80–87. IEEE, 2009.

[6] P. Henry, M. Krainin, E. Herbst, X. Ren, and D. Fox. Rgb-d mapping: Using depth cameras for dense 3d modeling of indoor environments. In *In the 12th International Symposium on Experimental Robotics (ISER*. Citeseer, 2010.

[7] P. Henry, M. Krainin, E. Herbst, X. Ren, and D. Fox. Rgb-d mapping: Using kinect-style depth cameras for dense 3d modeling of indoor environments. *The International Journal of Robotics Research*, 31(5):647–663, 2012.

[8] S. Hong, J. Jung, S. Kim, H. Cho, J. Lee, and J. Heo. Semi-automated approach to indoor mapping for 3d as-built building information modeling. *Computers, Environment and Urban Systems*, 51:34–46, 2015.

[9] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004.

[10] R. Martin-Brualla, Y. He, B. C. Russell, and S. M. Seitz. The 3d jigsaw puzzle: Mapping large indoor spaces. In *Computer Vision–ECCV 2014*, pages 1–16. Springer, 2014.

[11] N. Snavely and S. M. Seitz and R. Szeliski. Photo Tourism: Exploring Photo Collections in 3D. In *ACM Transactions on Graphics*, volume 25(3), pages 835–846, July 2006.

[12] M. Noreikis et al. Image based indoor navigation. 2014.

[13] S. Agarwal and Y. Furukawa and N. Snavely, et al. Building Rome in a Day. In *Commun. ACM*, Oct 2011.

[14] N. Snavely, S. M. Seitz, and R. Szeliski. Modeling the world from internet photo collections. *International Journal of Computer Vision*, 80(2):189–210, 2008.

[15] N. Snavely, S. M. Seitz, and R. Szeliski. Skeletal graphs for efficient structure from motion. In *CVPR*, volume 1, page 2, 2008.

[16] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon. Bundle adjustment a modern synthesis. In *Vision algorithms: theory and practice*, pages 298–372. Springer, 2000.

[17] M.-D. Yang, C.-F. Chao, K.-S. Huang, L.-Y. Lu, and Y.-P. Chen. Image-based 3d scene reconstruction and exploration in augmented reality. *Automation in Construction*, 33:48–60, 2013.

A survey on communication protocols and standards for the IoT

Gayathri Srinivaasan
Student number: 467122
gayathri.srinivaasan@aalto.fi

Abstract

The Internet of Things(IoT) is a novel paradigm which defines communication between a variety of products and devices in the real world, via objects such as Radio-Frequency IDentifaction (RFID) tags, sensors, actuators, mobile phones, etc., by way of unique addressing schemes [1]. IoT will have tremendous impact on everyday-lives of people in a variety of fields such as automation, industrial manfucturing, logistics and e-health. However, many challenges must be overcome before IoT is fully realized, the major issue being the interoperability between the devices that form the IoT. Despite many standardization requirements established, there is still lack of generic and standardized interfaces for definitive communication between the IoT devices.This paper surveys some of the communication protocols and assesses them against the requirements for a generic, application-level communication interface.
This paper is divided into 4 sections. Section 1 discusses the significance of the IoT and the need for interoperability between the IoT devices. Section 2 lists the functional requirements established by the Open Group, that need to be satisfied by any IoT messaging standard and also describes some of the existing messaging standards. Section 3 compares the messaging standards and their specifications described in Section 2 against the functional requirements established by the Open Group.

KEYWORDS: Internet of Things, Communication standards and protocols, RFID, Interoperability, Standardization requirements, Open Group, O-MI, O-DF, oBIX, CoAP, XMPP, MQTT.

1 Introduction

The Word Wide Web connects millions of people all over the world at any point of time. The modern internet era targets a world where physical objects and beings may efficiently interact with each other as and when needed. With the Internet of Things, communications are not limited to people-to-people or people-to-computers, but rather between people and objects and also between objects [3]. The IoT is not a single technology, rather an advanced paradigm in which the things/objects/devices/users are connected to each other. For instance, the street lights being networked and entities such as embedded sensors, augmented reality, near field communication being integrated into decision support, etc. are examples of the IoT in everyday life [8]. This in turn creates many business opportunities and adds to the complexity of the IT domain. With the growing number of Business to Business (B2B) infrastructures, there is a need for seamless communication between the complex business procedures at minimized ICT costs, while providing better services to users. The key challenge is the way the communications and the control techniques have evolved differently across the heterogenous business sectors. To accommodate this diversity, there is a need to address the interoperability requirements between the IoT applications. Before implementing the IoT, any organization has to make sure that their ICT infrastructure is flexible enough to establish information flows between various kinds of devices, products and information systems. The need for standardized and flexible communication technologies will create potential for several application areas in the future. Various standardization activities are currently in progress in the scientific community. This paper surveys some of the prominent standardization activities and assesses them against the requirements for a generic, application-level communication interface.

2 Functional requirements of the IoT

The basic requirement for the IoT infrastructure to emerge successful is the interoperability between heterogenous communication interfaces. The Open Data Format (O-DF) and the Open Messaging Interface (O-MI) standards (formerly known as Quantum LifeCycle Management standard) emerged out of the PROMISE EU FP6 project, where real-life applications required the collection and management of product information for many domains involving heavy and personal vehicles, household equipment, phone switches, etc. [3]. Based on the needs of those real-life applications, the requirements listed in Fig 1 were identified by the Open Group. These requirements aim to provide generic and standardized application-level interfaces to enable any product or device to exchange information as flexibly as possible for a robust solution [3].

This section introduces O-MI/O-DF messaging standards and their properties. It also details four other messaging protocols such as oBIX, CoAP, XMPP, MQTT and makes a comparison of the protocols with the O-MI/O-DF messaging standards by assessing them against the functional requirements of the Open Group as listed in Fig 1 [3] .

2.1 IoT Standardization Landscape:

The Internet of Things (IoT) paradigm has a wide scope and the standards landscape is infact large and complex. For successful interoperability between various domains, the

- Possible to implement for any kind of instances as independently of the application domain as possible
- Possible to implement for any kind of information systems, including embedded and mobile systems
- Support for “synchronous” messaging such as immediate read and write operations, including “client-poll” subscriptions
- Not restricted to one communication protocol only, it must be possible to send messages using protocols such as plain HTTP, SOAP, SMTP, as file copies, etc.
- Possibility to create ad hoc, loosely coupled, time-limited information flows “on the fly”
- Peer-to-peer communication possibility for all devices, i.e., client and server functionality can be implemented for any device, depending on available processing power, network connectivity, etc.
- Handling mobility and intermittent network connectivity, i.e., support for asynchronous messaging capabilities that imply for instance message persistence, time-to-live, etc.
- Context-dependent discovery of instances, instance-related services and meta-data about them
- Support for context- and domain-specific ontologies
- Queries by regular expressions for retrieving information about more than one instance and more than one kind of information
- Historical queries, i.e., retrieving values between two points in time

Figure 1: Functional Requirements of IoT messaging

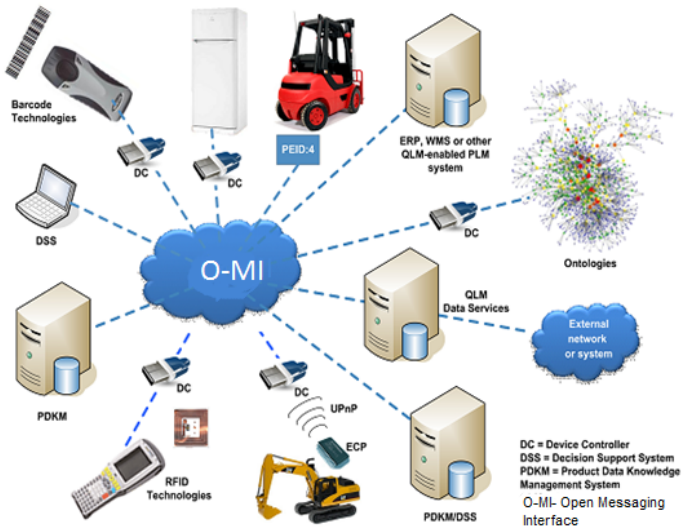

Figure 2: O-MI architecture [3]

standards are highly critical to ensure co-operation between the domains and enable the realization of proper "Internet of Things".
The Internet of Things European Research Cluster (IERC) is working to create a reference for pre-standardization activities for the European Commission IoT research projects [8]. There are various standardization activities currently in progress and being proposed. This section gives an overview of 5 of the messaging standards/protocols proposed by international standard organizations, including OpenGroup, CEN/ISO, IETF, OASIS and IEEE.
The basis of IoT is the web and some of the established standards of the IoT are the HTTP(S), FTP and SMTP. But they are not suitable for low-power, low-memory, processing constrained devices [3].Various organizations have proposed standards namely XMPP, MQTT, CoAP which are aimed at resource-limited devices and run directly on TCP and/or UDP.

The following sections describe the various standardization activities established or in progress and assess them against the IoT requirements by the Open Group.

2.1.1 O-MI and O-DF standards

The O-MI/O-DF standards by the Open Group aim to allow communication between intelligent entities to enable exchange of IoT information in ad hoc, loosely coupled ways. These standards combine the main features of asynchronous, enterprise messaging protocols with that of instant messaging protocols to enable peer-to-peer communications [3]. While the web uses the HTTP protocol to transmit information in HTML format, the O-MI transports messages represented using the O-DF. In order to describe complex data structures as flexibly as possible, both the O-MI and O-DF specifications are written using XML schema.

The O-MI standard satisfies most of the functional requirements established by the Open Group as given in Fig 1.

- **Applicability:** The O-MI messaging standards can be implemented in any kind of information systems, including embedded systems and mobile devices.
- **Protocol Binding:** The O-MI standards are highly non-dependent on any specific communication protocol. O-MI is protocol agnostic which means that the O-MI messages can be exchanged using HTTP, SMTP, MQTT, CoAP, etc. [3].
- **Synchronicity and Client-Server model:** The O-MI standard allows synchronous (real-time) communication between devices. They support immediate read-write operations, including client-poll subscriptions.
- **Subscription:** The O-MI messaging standard follows Observer subscription model where a device/node can add itself as an observer of events that occur at another O-MI node. This is significantly different from the publish-subscribe mode which assumes the use of a "high-availability server". Hence Observer model is more suitable for IoT applications where products might communicate with each other directly [5].
- **History query:** The O-MI/O-DF support historical trends of data and allow querying of history data between any two points in time.
- **Publication and Discovery:** The O-MI/O-DF enable publication and discovery of instances, instance-related services with RESTful URL-based queries. The O-MI/O-DF also support mobility and intermittent connectivity by specifying time-to-live.

2.1.2 oBIX

oBIX (Open Building Information Xchange) is an industry-wide initiative by the OASIS Group, to define XML and web services-based standard for communication between building and electrical systems, and enterprise applications [7].

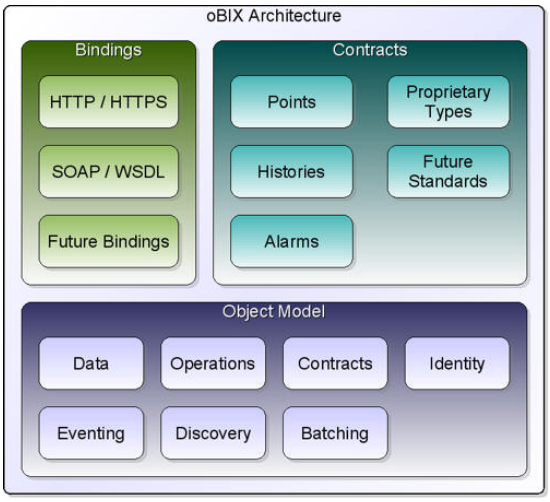

Figure 3: oBIX architecture [7]

oBIX is a standard web services protocol that provides the mechanical and electrical control systems, easy access to the status of their operations, performance, problems or faults which require analysis or attention. The most important feature of oBIX is extensibility through a concept called **'contracts'**. A contract is a list of all the patterns a complex piece of data conforms to [4]. In simple terms, contracts in oBIX are synonymous to classes in object oriented systems. The oBIX specification is majorly defined through contracts. The significance of contracts is that a new contract can be added without changing the existing oBIX schema[7] .

oBIX satisfies some of the functional requirements established by the Open Group as given in Fig 1.

- **Applicability:** oBIX is designed to provide access to the embedded software systems which sense and control the world around us, via simple XML . It is specifically used for building and mechanical systems for diverse M2M communication.
- **Protocol Binding:** oBIX binds with SOAP protocol to interoperate with web services and also with HTTP to become a RESTful standard.
- **Synchronicity:** oBIX, much like the WWW, is a big web of XML object documents hyperlinked together using URIs [7]. oBIX supports 3 request types:-
  - **Read:** it returns the current state of an object
  - **Write:** it updates the state of an object
  - **Invoke:** it triggers an operation on an object

  In addition to asynchronous 'Read' and 'Write' operations, oBIX also supports the 'Alarm' feature which indicates a condition that requires notification of either a user or another application [5]. For example, safety alarms in buildings on the occurence of an event.
- **Client-Server model:** The oBIX architecture representated in Fig 3 is based on a client/server network model [4].
  - **Server:** The software with oBIX enabled data and services respond to requests from clients over a network.
  - **Client:** The software which makes requests to the servers over a network to access oBIX enabled data and services.
- **History query:**

  oBIX supports historical trends of data by defining a list of time stamped values. oBIX allows query and roll up of history data [7]
  - **History Object:** a normalized representation of the history
  - **History Record:** a record of a point sampling at a specific timestamp
  - **History Query:** a way to query history data
  - **History Rollup:** a mechanism to summarize an interval of time in the history data.

To summarize, the oBIX protocol is a suitable communication protocol designed to provide access to a wide range of devices and standardize communication between them.

2.1.3 CoAP

CoAP is the Constrained Application Protocol from the CoRE (Constrained Resource Environments) IETF group. It is specifically designed for use with low-power, resource constrained nodes and lossy networks. Unlike HTTP-based protocols, CoAP operates over UDP to save bandwidth in resource constrained devices. CoAP is based on the REST architecture. In order to overcome the problem of resource constraints, CoAP optimizes the length of the datagram and provides reliable communication. The protocol is particularly intended for small low power sensors, switches, valves and similar components that need to be controlled or supervised remotely via Internet [10].

Fig 4 portrays the main features of the CoAP protocol.

CoAP satisfies some of the functional requirements established by the Open Group as given in Fig 1.

- **Applicability:**
  Information appliances, control equipment and communication equipment in smart home networks have the characteristics of low-cost and lightweight. Thus, CoAP could be viewed as the best protocol choice for home communication networks or any domain with the need for communication with resource constrained devices [2].
- **Subscription:**
  The state of a resource on a CoAP server can change over time. Since continuous polling increases the complexity of resource constrained devices, CoAP supports the 'observe' feature. The client sends an 'observe' request to a resource on the CoAP server. From then on,

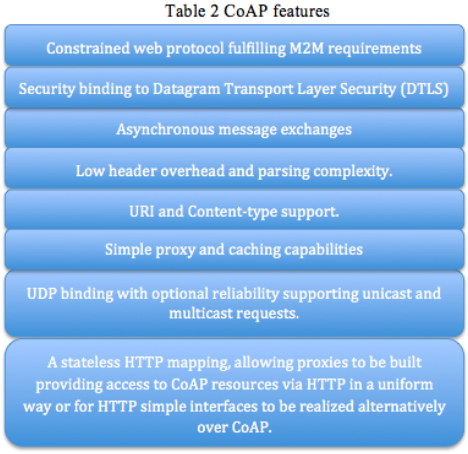

Figure 4: CoAP features [2]

the CoAP server notifies the client of any change to the current state of the resource and returns the resource representation [2].

- **Synchronicity & Client-Server model:**
  Like HTTP, CoAP supports the client/server model for communication. CoAP implementation allows peer to peer communication with dual roles of the client and the server. A CoAP client sends a request for an action on a resource (identified by the URI) in the server. The server then sends a response with a response code; this response may include a resource representation. Unlike HTTP, CoAP handles these message exchanges asynchronously using UDP [10].

- **Protocol Binding:**
  CoAP is designed to interoperate with HTTP and the web through simple proxies. Because CoAP is datagram based, it may be used on top of SMS or other packet based communications protocols.

- **Resource discovery:**
  CoAP provides inbuilt support for content negotiation and resource discovery allowing devices to query each other to find ways of exchanging data [2].

Thus, CoAP is best suited for resource constrained devices for light-weight communication and overcomes the complexitiy involved with HTTP.

2.1.4 MQTT

MQTT (formerly Message Queue Telemetry Transport) is a lightweight, broker-based publish/subscribe messaging protocol designed to be simple and easy to implement. It has recently been established as a standard by the OASIS technical committee. MQTT is aimed at constrained environments such as high latency, low bandwidth, unreliable and high cost networks where the client applications have limited processing capability [6].

MQTT fulfils the following functional requirements as put forward by the Open Group in Fig 1.

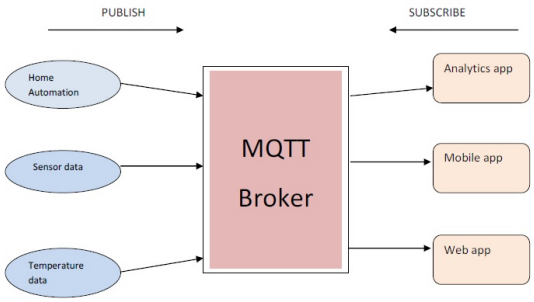

Figure 5: MQTT Publish-Subscribe model

- **Applicability:**
  MQTT is applicable to resource constrained devices. Unlike HTTP, MQTT is agnostic of the data content and deploys a Publish-Subscribe messaging protocol, which is highly efficient when compared to multiple HTTP GET and POST request model. This is particularly useful for lightweight M2M communications [6].

- **Subscription & Client Server model:**
  MQTT operates over TCP and the publish/subscribe message pattern of MQTT provides one-to-many message distribution and decoupling of applications. The clients may subscribe to multiple topics and the messages are published as a specific topic name (PUBLISH). The clients open TCP connections with the MQTT broker which handles routing of messages. Every client subscribed to a topic receives every message published to that topic [4]. Fig 5 shows the architecture of the MQTT protocol.

- **Synchronicity:**
  With the publish- subscription model, MQTT supports asynchronous communication between the server and the clients by providing one-to-many message distribution and decoupling of applications. When publishing messages,the clients may request persistence of the message by the broker [6]. One significant example of the MQTT asynchronous communication is the Facebook Mobile Messenger.

- **Protocol Binding:**
  The MQTT protocol runs over TCP/IP, or over other network protocols that provide ordered, lossless, bidirectional connections [6].

Hence, MQTT is best applicable to networks which are expensive and unreliable and possess limited processing capabilities.

2.1.5 XMPP

Extensible Messaging and Presence Protocol (XMPP) is a communications protocol for real-time communication using XML. It provides a way to send small pieces of XML

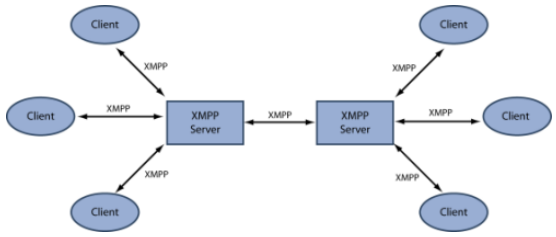

Figure 6: XMPP Architecture [9]

from one entity to another in real time. The protocol was originally developed by the Jabber open-source community and formalized later by the XMPP working group formed by the Internet Engineering Task Force (IETF). XMPP is implemented using distributed client-server architecture and enables asynchronous, end-to-end exchange of structured data by means of direct, persistent XML streams [5]. XMPP was originally developed for instant messaging to connect people via text messages. XMPP satisfies some of the functional requirements established by the Open Group as given in Fig 1.

- **Applicability:**
  XMPP addresses a device using the addressing scheme "name@domain.com". In the context of IoT, this addressing scheme is advantegeous to communicate to a remote device via the web. Some of the significant features of XMPP are - Channel encryption, Authentication, Presence, One-to-one messaging, Notifications, Multiparty messaging, Service discovery etc. XMPP provides a great way, for instance, to connect your home refrigerator to a web server and access it via mobile phone.
- **Client-Server model:**
  XMPP follows client-server architecture, as displayed in Fig 6, where a client utilizing XMPP accesses a server over a TCP connection, and servers also communicate with each other over TCP connections.
- **Subscription:** XMPP includes the ability for an IoT enabled device/entity to advertise its network availability known as "presence" to other entities. This is performed in the form of a Publish-Subscribe model [9].
- **Synchronicity:** In HTTP, the client sends a request to a server and then waits for a reply before it makes another request. By contrast, in XMPP the client can 'pipeline' requests to the server or to other entities and then receive replies as they appear. The occurrence of certain events also triggers information that is pushed to the client.
- **Resource discovery:**
  XMPP supports 'Service discovery' feature which enables one entity to discover the features supported by another entity. For instance, find information about chat rooms in a chat service.

Thus, XMPP is a real-time communication protocol best suited for a distributed client-server architecture.

Functional Requirement	QLM	oBIX	CoAP	MQTT	XMPP
Applicability	x	x	x	x	x
Protocol Binding	x	x	-	x	-
Synchronicity	x	-	x	-	-
Subscription	x	x	x	x	x
History Query	x	x	-	-	-
Client-Server model	x	x	x	x	x
Resource discovery	x	-	x	-	x

Table 1: Comparison of protocol specifications

3 Comparison of the messaging protocols

In addition to the O-MI/O-DF messaging standards, Section 2 assessed the properties of four IoT messaging standards against the functional requirements established by the Open Group. It is clear that the O-MI/O-DF standards satisfy most of the functional requirements established by the Open Group, when compared to the other protocols. It is observed that protocols such as CoAP, MQTT are suited for communication with low-power, low-resource devices. However, it is also important to note that the O-MI/O-DF standards are not limited in usage with low-power, resource-constrained devices and can be used with a myriad of devices. The comparison shows that the O-MI standards satisfy majority of the functional requirements which are necessary for interoperability between business infrastructures in order to provide a robust IoT solution. The Table 1 summarizes the results of the comparison in Section 2 . The 'x' marks denote that the particular requirement is satisfied by the protocol.

4 Conclusion & Future Work

This paper surveyed four different messaging protocols namely oBIX, CoAP, MQTT and XMPP in addition to the O-MI messaging protocol and assessed their features against the functional requirements put forward by the Open Group. The paper discussed the key features of these protocols that are required by any business solution to exchange information with one another. It is clear that some of the recent messaging protocols are targeted at low-power and resource-constrained devices and employ techniques to bypass HTTP standard to enable device communication. It is also obvious from the comparsion that the O-MI standards are more flexible and satisfy majority of the functional requirements required for device interoperability. O-MI and other protocols discussed are an essential step to enhance product lifecycle management and enable exchange of product information in order to provide the right service at the right time, anywhere and to anyone. Further research could be carried out to identify more standards to enable complete interoperability in the IoT and achieve successful collaboration of diverse business infrastructures.

References

[1] L. Atzori, A. Iera, and G. Morabito. The internet of things: A survey. *Computer Networks*, 54(15):2787 – 2805, 2010.

[2] X. Chen. Constrained Application Protocol for Internet of Things. `http://www.cse.wustl.edu/~jain/cse574-14/ftp/coap/index.html`. [Accessed Apr 18, 2015].

[3] K. Framling, S. Kubler, and A. Buda. Universal messaging standards for the iot from a lifecycle management perspective. *Internet of Things Journal, IEEE*, 1(4):319–327, Aug 2014.

[4] A. Hansen. Open Building Information Exchange- oBIX Unbound. `http://www.automatedbuildings.com/news/may06/articles/obix/060425112552obix.htm`, 2006. [Accessed Apr 18, 2015].

[5] S. Kubler, M. Madhikermi, A. Buda, and K. Framling. Qlm messaging standards: Introduction and comparison with existing messaging protocols. In T. Borangiu, D. Trentesaux, and A. Thomas, editors, *Service Orientation in Holonic and Multi-Agent Manufacturing and Robotics*, volume 544 of *Studies in Computational Intelligence*, pages 237–256. Springer International Publishing, 2014.

[6] MQTT. MQ Telemetry Transport. `http://mqtt.org/documentation`. [Accessed Apr 18, 2015].

[7] OASIS. Open Building Information Exchange Technical Specification . `https://www.oasis-open.org`, 2014. [Accessed Apr 18, 2015].

[8] O.Vermesan and P.Friess. *Internet of Things - From Research and Innovation to Market Deployment*. River Publishers, 2014.

[9] P. Saint-Andre. Extensible Messaging and Presence Protocol (XMPP): Core. `http://xmpp.org/rfcs/rfc6120.html`, 2011. [Accessed Apr 18, 2015].

[10] C. B. Zach Shelby, K. Hartke. Constrained Application Protocol (CoAP). `http://tools.ietf.org/html/rfc72528`, 2014. [Accessed Apr 18, 2015].

Website reputation and classification systems

Sridhar Sundarraman
Student number: 467216
sridhar.sundarraman@aalto.fi

Abstract

The purpose of a website reputation rating system is to present to the user how safe a website is. The safety rating is determined by taking various factors into consideration, including but not limited to, likelihood of phishing or scamming, presence of viruses or malware, hosting of offensive or illegal content and misuse of private user data. Existing systems use manual expert evaluation, crowd-sourcing to obtain input ratings from users, and also machine learning techniques to determine the rating. This paper presents a survey of the recent trends in assessing website reputation ratings from a research perspective and the techniques used by commercially available systems. In addition, it presents a commentary of the significance and comparison of different approaches employed in contemporary website safety rating determination techniques.

KEYWORDS: Website Safety, Reputation System, Website classification, Safety Features

1 Introduction

With the number of malware and rootkit threats still growing at a significant rate [12], it has now become essential to employ defensive mechanisms to protect users from the recent threats. Web reputation is one such method to protect against malicious content on the web. A website is considered malicious if its intention is to disrupt normal user operation or to gather sensitive data or to gain access to systems without an user's consent. These malicious websites are classified under the umbrella term of "Malware" and may include any type of website that introduces viruses, worms or trojans to the user's systems and also includes spywares and adwares. Web reputation systems involve a comprehensive security assessment of a website based on its potential to be a threat and the providing of a representation of the safety of the website in the form of a rating, thus determining the risk factor of the website. A common way among web reputation systems to provide such a rating is in the form of a score from 0 to 100, where the lower the score, the riskier the website is. Table 1 represents a typical way to broadly classify websites based on their safety rating [17]:

- High Risk (1-20)
- Suspicious (21-40)
- Moderate Risk (41-60)
- Low Risk (61-80)
- Trustworthy (81-100)

Anti-virus products have different response types to new zero-hour and zero-day threats, which are attacks that take advantage of the fact that there has been no time for a developer to provide a fix/patch for the vulnerability. For such cases, web reputation systems act as additional defence mechanisms by enhancing the security coverage. In some particular cases where the threat has not yet been detected, these systems may be the only level of protection available to users. Some of the features of these systems include their ability to predict the security of a website even before an actual threat is detected. This is achieved by keeping previous records of infected sites and using this data to determine the reputation score. In web filtering, web reputation can also be combined with content categorization to identify the type of a website, e.g. "News and Media", "E-Commerce", and "Social". The categories may also include information about the security risks involved, such as, "Phishing" or "Malware".

Web reputation scores may be used in different scenarios. Known compromised sites which are categorized as being "High Risk" band may be blocked. For websites of moderate risk, a warning may be presented to the user informing her of possible risks, thus providing proactive protection.

Section 2 lists some factors that are typically analysed in websites to compute the safety ratings. Section 3 to 5 present the various processes involved in a typical website reputation system, with emphasis on recent research trends. These include the process in which data is gathered about the known safety of websites, the way in which safety ratings are presented to users and the properties that are analyzed to determine the safety. In addition, section 7 presents various commercially available website reputation systems and the approach they use in their prediction of website safety.

2 Safety variables

Typical web reputation systems use a large number of factors and a machine learning technique to predict the safety rating of a website. Instead of using a manual process, a predictive model is built using a known dataset of millions of examples of websites whose safety rating is known by other means as the training dataset. This model is then able to assess the rating of any arbitrary website whose rating is unknown.

The sample input provided to the classifiers can be obtained by analysing features of the websites such as location of website, ISP/Webhost, IP neighbourhood [17], structural features from HTML and Javascript embedded in the site [9], content based feature set from malicious outbound links and

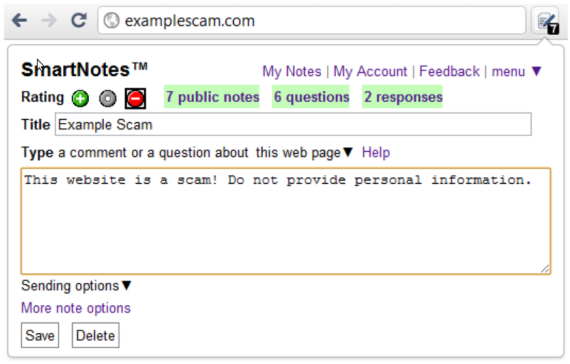

Figure 1: Crowdsourced rating collection by SmartNotes [10]

Figure 2: Groupsourced rating collection by FAR [13]

textual content [7]. Static features that could be extracted from the website may include [9], for example, presence of iframe tags in HTML, occurrence of hidden UI elements, the presence of meta refresh tags and embed tags.

When legitimacy of the website is taken into consideration, the factors may include the domain age and the number of IP addresses for the site. Unless it is clear that a website is legitimate, we should treat new websites as suspicious by default. Many web reputation systems also maintain known threat history for websites based on the assumption that a website which has been infected in the past is more likely to be risky than a website with a clean past threat record.

3 Data Gathering

3.1 Crowdsourcing

Crowdsourcing, as already noted above, is widely used in gathering safety ratings from user. These ratings, in the simplest case, are aggregated and presented to users in a meaningful way. Crowdsourcing is analogous to user-review systems that are commonly implemented in E-commerce platforms, where users are asked to share their experiences of particular products that they have purchased.

Web of Trust [6] (discussed in detail in section 7.3) is a commercial system that uses this approach to present ratings to the users in the form of a browser plugin. Fink et. al., [10] propose a crowdsourcing architecture where they users are asked to share their experiences about web threats. Their platform "SmartNotes" allows a user to rate websites, post comments and also to ask and answer questions related to safety of websites. In addition to integrating their system with existing crowdsourced systems used to answering questions, they also employ machine learning and natural language processing to analyze the user feedback.

3.2 Groupsourcing

Groupsourcing is an approach that involves gathering data from a user's social circles instead of the public in general. This approach has the advantage of the fact that individuals within a social group can be trusted more, in a social sense, than a random stranger. Lin [13] uses groupsourcing to identify unsafe content in online social networks. They build a predictive model based on the groupsourced dataset and apply Support Vector Machine (SVM) as a classifier with which they predict the rating level for a given website.

3.3 Advantages and Disadvantages

In many systems that use machine learning techniques, the dataset collected using crowdsourcing is used as the ground truth for building a predictive model. The ratings provided by crowdsourcing systems usually can be considered to be relatively accurate, as intentional manipulation of ratings by some malicious users is usually offset by the ratings provided by a multitude of users. Crowdsourcing is relatively cheaper than conducting user studies and provides a wide diversity of input.

The main disadvantage of crowdsourcing is time lag, which is a time delay between the launch of a website and the time when a rating is available for that website. This is because the task of outsourcing a job to a large group of people is inherently time consuming. Groupsourcing partially tackles this problem, as ratings available from just a single user within the social circle is motivation enough to conclude the rating, whereas crowdsourcing needs sufficient number of users to perform aggregation of results. Another disadvantage is that crowdsourcing is not scalable to the entire Internet. The coverage of ratings that can be obtained is only minimal. [7] provide a representation of the number of websites that are rated by Web of Trust [6] among the most popular one million websites as collected by Alexa (http://www.alexa.com).

4 Representation of safety

4.1 Blacklists and warnings

Blacklisting is the most common approach to fighting scams and malware distribution in the Internet. A database of malicious websites is maintained and provided to others by several online services including Google, Symantec and Securi. It is common among browsers to display a warning if an user tries to visit a website which is found in any of the blacklists. Malicious websites that are added to the blacklists are

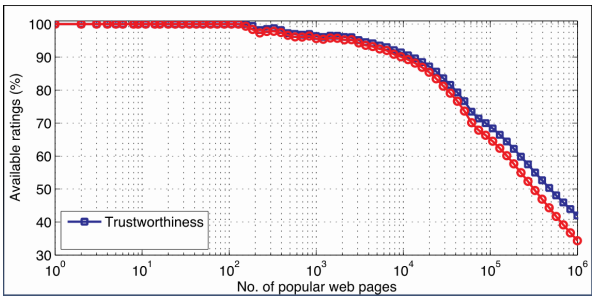

Figure 3: Availability of ratings from WOT for most popular 1 million websites as collected by Alexa [6, 7]

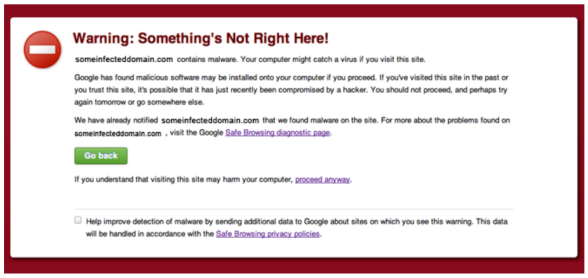

Figure 4: Warning shown by Google Chrome [2]

usually from abuse reports from users and other authorities such as Google, Bing, Norton Safe Web and McAfee SiteAdvisor . These websites may have been reported to have distributed malware, to have suspicious activity as revealed by inconsistencies from search engines or pose Phishing attach threats. Fig 4[2] shows the warning displayed by web browsers Google Chrome when a user visits a blacklisted website. Other popular websites such as Mozilla Firefox and Apple Safari also show similar warnings before letting the user visit.

Although blacklists may be useful to protect users from infamous malicious sites, they have several limitations. They may not contain very recent scam sites, as well as sites that may have moved to a new domain. Moreover, intentionally biased or inaccurate reports may result in legitimate websites being added to blacklists. They also should be updated periodically lest they becoming obsolete.

Apart from warnings, as noted in section 1 the safety can be represented as a score [6], or as categories of varying level of riskiness [4, 3]. Table 1 shows the way difference reputation systems present the safety of a website.

Reputation System	Safety Representation
WOT	Trustworthiness Score: 0-100 Child Safety Score: 0-100
McAfee SiteAdvisor	Color-coded glyphs: Red: Website is safe Yellow: Some issues with the site Red: Serious issues with the site Gray: No rating available Blue: Internal site or private IP range Black: Phishing site
Norton Safeweb	Color-coded glyphs: 
F-Secure Search	Categories: Safe, Allowed, Suspicious, Harmful, Unknown and Denied

Table 1: Safety representations

5 Properties used in safety determination

5.1 Suspicious URLs

In order to tackle the limitations of blacklists, Ma et. al., [14] propose an approach for classification of websites using machine learning classifiers based on lexical features and host-based features extracted from the website. They consider the task of URL reputation prediction as a binary classification problem, with malicious websites being positive samples and benign URLs as negative samples. The lexical features that they consider include properties such as the length of the host name, length of the URL in its entirety, and the number of dots in the URL. In addition, they consider binary features for each token in the URL (parts delimited by a '.'). The host based features considered include IP address properties, WHOIS properties, domain name properties and geographic properties. They claim that the host based features would be sufficient to describe the host of the malicious sites, the owner of the site, and the way in which the site is managed. These features are used to build models which are then tested by classification algorithms such as Naive Bayes, SVM and logical regression. In their evaluation, they report a 14.8% false positive rate (FPR) and 8.9% false negative rate (FNR).

Ma et. al., do not consider any other features other than above-mentioned lexical and host-based features. Particularly, they argue that extracting features from the website's content for classification is a slow process and may be resource intensive. Although, it is possible that malicious websites serve different contents to clients based on IP addresses, the contents is still a significant factor that determines the safety of the website. While classification using features based only on the URLs may be applicable in any context, there is a possibility of malicious websites cloaking URLs using shortener services or client side redirection directives.

5.2 DNS and Web server relationships

Many attackers employ a centralized exploit server from which malicious content is served to many websites that the attacker has control of. Involving many websites, in such a manner, renders investigation of the attack difficult. Seifert et. al., [16], present a method to identify malicious websites by identifying attributes that characterize the servers referenced by the web page. If the servers involved are malicious in nature, then this affects the safety rating of the website

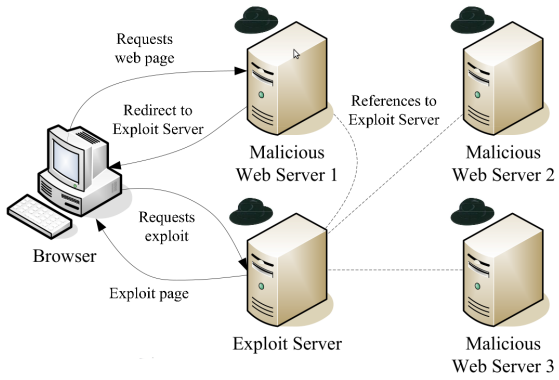

Figure 5: Server relationship structure in a typical attack [16]

in question. Their motivation to propose such a system is to tackle limitations of client honeypots that are commonly used in detecting servers that launch drive-by-download attacks. These honeypots are resource intensive and have a high tendency for predicting false negatives. Moreover, since typical honeypots make use of a vulnerable client to track unauthorized changes during interaction with a malicious server to detect the maliciousness, if there is no interaction, the site may go undetected.

The approach proposed by Seifert et. al. involves analyzing relationship between servers. Two main types of servers involved are the DNS servers and web servers. In a typical attack, the common structure of inter-related servers include a centralized exploit server. Here, the exploit server serves malicious content via several web servers which it controls. The actual serving of the exploit pages can be done via iframes and redirects and exploitation kits exist that support this structure. Fig 5 shows a typical structure of servers involved in an attack. In order to collect training dataset, they have employed a client honeypot to monitor and capture network traffic between vulnerable clients and malicious server and in this process have identified what DNS servers are involved. This data is then fed into a J4.8 decision tree learning algorithm. This predictive model can assess whether a web page is malicious or not.

This approach has reported a high rate of FN (25.5%) but a low rate of FP (2.6%). Also, this approach should be periodically updated to counter the attackers who change their technique.

5.3 Passive Domain Analysis

Employing botnets is a common attack scenario where a multitude of end-user systems are compromised and taken over by attackers. These clients are then turned into bots which are then used to launch Distributed Denial of Service (DDoS) attacks, steal sensitive data, and also for spamming. In order to setup such an infrastructure, the clients must report back to the exploit server, which requires that the IP address of the server be hardcoded on the client side. This results in a single point of failure for the attackers, where if the IP address is investigated and found, it may bring down the entire botnet. To avoid such a case, attackers make efficient use of Domain Name System (DNS). DNS gives the flexibility of changing the IP addresses dynamically and also to hide behind proxy servers, making it difficult to trace back.

Bilge et. al., [8] demonstrate that, by analysing DNS traffic and by studying the behaviour of known malicious and benign domains, we can identify features that help determine the maliciousness of a domain. The features that they have identified are classified as Time-based features, DNS answer-based features, TTL Value-based features and Domain Name-based features. In their system called EXPOSURE, they have passively monitored DNS traffic from Security Information Exchange (SIE). EXPOSURE is trained using a dataset of over 100 billion DNS queries and 4.8 million domain names and has predicted over 3000 malicious domain names when tested within a commercial Internet Service Provider (ISP). Their classifier uses the J48 decision tree algorithm. When tested with a recorded data set, EXPOSURE reported that 5.9% of domains out of 300,000 domains were malicious. When cross-checked with data from McAfee and Norton, a false positive rate of 7.9% is estimated.

As with any website reputation system that extracts features as inputs for classifiers, EXPOSURE suffers from the limitation that attackers may manipulate these features that are deemed to report maliciousness. Again, as with any other system that uses machine learning techniques, the efficiency of the system depends on the training set.

6 Comparison and Discussion

Many of the approaches, as noted above, take the static/structural features of a website as the major safety variable. These features are chosen based on assumptions that they may have a correlation with the safety of a website. For example, the approach by [14] takes only the URL and features derived from the URL as safety variables. Considering dynamic features and content-based features (using topic modelling techniques for extraction) may also provide better accuracy. These, along with empirical cumulative distribution of embedded function (ECDF) of the ratings of the embedded links in a website, are being considered in an active research carried out by the Secure Systems research group at Aalto University [7]. Analyzing DNS and web server relationships to determine safety [16] may be applicable only to a specific subset of attacks, particularly, the attack scenario mentioned in section 5.2 and is not generalizable to cover a wider range of threats.

The results of any machine learning algorithm is expected to have as less a FP and FN rate as possible and these rates have a direct correlation with its accuracy. A high FN rate implies that the system has a high probability to predict a malicious website as beingn and a high FP rate implies that it has a high probability to predict a benign website as malicious. Out of the approaches described in the previous section, [14] and [8] have similar FN and FP rates, but [16] has a very high FN rate.

Many of the reputation systems consider the prediction of the safety as a classification and so the type of the classifier used also has an impact on the accuracy of the classi-

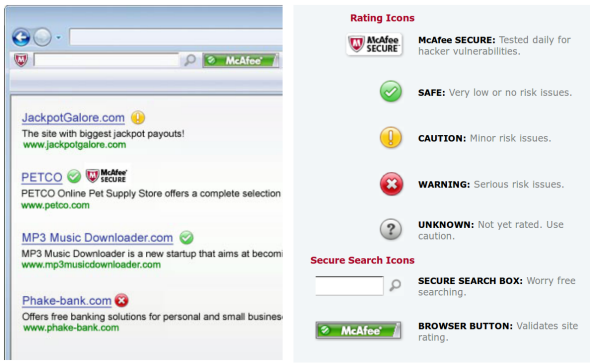

Figure 6: McAfee SiteAdvisor and their rating categories [4]

fication. For example, [14] have experimented with Naive Bayes, SVM and logical regression; [16, 8] have both used J.48 decision tree algorithm, which is a an open source Java implementation of the C4.5 algorithm [15].

7 Commercial Website Rating Systems

Many anti-virus companies maintain a website reputation system and make it available to the end-user via a browser add-on/plug-in which then provides visual color-coded glyphs besides search results from search engines like Google and Bing. The approaches taken by some of the commercially available web reputation systems are discussed below.

7.1 McAfee TrustedSource and SiteAdvisor

McAfee's TrustedSource [5] is an Internet reputation system that aids their SiteAdvisor software [4]. SiteAdvisor is installed as a browser plugin and shows site rating icons beside search results. They also provide an optional search box which filters search results automatically. These products alert the users to potential risks. The rating icons (as shown in Fig 6) comprise of four categories: "Safe", "Caution", "Warning" and "Unknown", which correspond to low risk, medium risk, high risk and unverified websites respectively.

Factors that SiteAdvisor analyze in websites include:

- Presense of downloadable files such as toolbars or screensavers, where there is possibility of them being bundled with viruses, spyware and adware. Websites earn a red rating, if there is any such presence of files.
- Possibility of spamware via e-mail subscriptions.
- Browser exploits or drive-by-downloads which may install trojans and keyloggers into the user's system.
- The reputation of the website, as reported by their TrustedSource system.

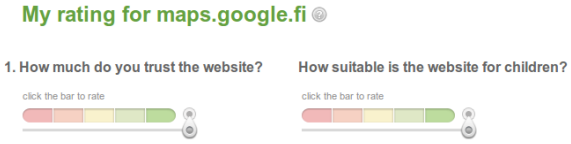

Figure 7: WOT's crowdsourcing approach: Collecting ratings from users [6]

- Vulnerabilities in E-commerce websites, where there are high possibilities of financial data theft.
- Annoyances like pop-ups and cookies.
- Outbound links from websites: Many websites, while not being malicious themselves, are setup just to trick users to visit other malicious websites.

7.2 Norton Safe Web

Norton's SafeWeb [3] is a reputation system that also provides users with a browser toolbar that alerts users to the safety ratings of websites. The ratings are, as is the case with SiteAdvisor, shown in 4 categories of varying risks. The major difference between McAfee's SiteAdvisor and SafeWeb is that the latter requires users to sign in with an account and also asks for the users' experiences of websites, as a crowd-sourcing measure. SafeWeb, in addition to the factors analyzed by SiteAdvisor, also analyzes unsolicited browser changes, suspicious browser changes, trackware, Cybersquatting and Pay-per-click sites.

7.3 Web of Trust

As already noted in section 3.1, Web of Trust (WOT) [6] is a reputation service that employs a unique crowdsourcing approach, where a global community of users rate and review websites based on their personal experiences. WOT, again, shows glyphs/indications in the form of traffic lights beside search results and social networking sites such as Facebook and Twitter. Clicking on the traffic lights shows a pop-up which provides detailed information about the reputation of the specific website. WOT provides ratings in two dimensions: trustworthiness and child safety as an integer in the range of 0 to 100. Fig 7 shows a page where users can rate their experience of a website.

8 Usability of website reputation systems

Social acceptance or the usability is also a critical factor to be considered when designing a website reputation systems. This may dictate how useful a system actually is to a user.

Sucuri [1] show that blacklists and warnings shown by browsers is very effective. They mention that a website which has been blacklisted by a search engine loses nearly 95% of its organic traffic. This effectively implies that users take blacklists very seriously.

User studies conducted by Karvonen et. al., [11] show that visually prominent parts of a reputation system usually gets prominent attention from users. This is employed effectively by many of the commercial systems, as noted in section 7. Visually prominent reputation information includes images, symbols and statistical visualizations. The most common representation of website safety is by color-coded glyphs and icons which correspond to the level of risk that the website poses. These visual information is also usually combined with detailed textual information and pointers to how the reputation is actually determined.

The form in which a reputation system is available to the end-user is a factor that affects the usability. Many systems provide browser-addons, which highlights the safety rating beside search results. Some reputation system clients are only available as part of a larger enterprise security product suite, and not available as stand-alone. Norton SafeWeb provides its own search engine and also a web interface in which users can query the reputation giving an URL as the input.

9 Conclusion

This paper briefly introduced the definition and purpose of a website reputation rating system. It also presented a survey of trends in approaches to estimate the safety of a website and a brief look at some specific research studies. In order to overcome the limitations of traditional defence mechanisms such as blacklisting, reputation systems typically use machine learning techniques in which they build a predictive model based on a training dataset. Various features of the websites are considered by different approaches as factors that may indicate their maliciousness. This paper also discussed the effectiveness and limitations of such features. Finally, the paper gave an overview of contemporary commercial reputation systems, the features that they analyze and also the form in which present their ratings.

References

[1] Google Blacklisted Website: Understanding and Removing A Google Blacklist. https://sucuri.net/website-security/google-blacklisted-my-website. Accessed: 2015-04-18.

[2] Google Blacklist Warning: Something is Not Right Here! http://blog.sucuri.net/2012/07/google-blacklist-warning-somethings-not-right-here.html. Accessed: 2015-03-21.

[3] Is this Website Safe | Web Security | Norton Safe Web. https://safeweb.norton.com/. Accessed: 2015-03-21.

[4] McAfee SiteAdvisor Software - Website Safety Ratings and Secure Search. http://www.siteadvisor.com/howitworks/index.html. Accessed: 2015-03-21.

[5] TrustedSource - Internet Reputation System. http://trustedsource.org/en/home. Accessed: 2015-03-21.

[6] Web of Trust. https://www.mywot.com/. Accessed: 2015-03-19.

[7] S. Bhattacharya, O. Huhta, and N. Asokan. Lookahead: Augmenting crowdsourced website reputation systems with predictive modeling.

[8] L. Bilge, E. Kirda, C. Kruegel, and M. Balduzzi. Exposure: Finding malicious domains using passive dns analysis. In *NDSS*, 2011.

[9] D. Canali, M. Cova, G. Vigna, and C. Kruegel. Prophiler: A fast filter for the large-scale detection of malicious web pages. In *Proceedings of the 20th International Conference on World Wide Web*, WWW '11, pages 197–206, New York, NY, USA, 2011. ACM.

[10] E. Fink, M. Sharifi, and J. G. Carbonell. Application of machine learning and crowdsourcing to detection of cybersecurity threats application of machine learning and crowdsourcing to detection of cybersecurity threats, 2011.

[11] K. Karvonen, S. Shibasaki, S. Nunes, P. Kaur, and O. Immonen. Visual nudges for enhancing the use and produce of reputation information. In *Proceedings of the ACM RecSys 2010 Workshop on User- Centric Evaluation of Recommender Systems and Their Interfaces*, 2010.

[12] M. Labs. McAfee Labs Threats Report June 2014. Whitepaper, McAfee. http://www.mcafee.com/mx/resources/reports/rp-quarterly-threat-q1-2014.pdf.

[13] J. Liu. How to Steer Users Away from Unsafe Content. Master's thesis, University of Helsinki, Helsinki, 2014.

[14] J. Ma, L. K. Saul, S. Savage, and G. M. Voelker. Beyond blacklists: Learning to detect malicious web sites from suspicious urls. In *Proceedingsof theSIGKDD Conference. Paris,France*, 2009.

[15] S. Salzberg. C4.5: Programs for machine learning by j. ross quinlan. morgan kaufmann publishers, inc., 1993. *Machine Learning*, 16(3):235–240, 1994.

[16] C. Seifert, I. Welch, P. Komisarczuk, C. Aval, and B. Endicott-Popovsky. Identification of malicious web pages through analysis of underlying dns and web server relationships. In *Local Computer Networks, 2008. LCN 2008. 33rd IEEE Conference on*, pages 935–941, Oct 2008.

[17] G. Whitepaper. How Web Reputation increases your online protection. Whitepaper, GFI. www.gfi.com/whitepapers/web-reputation-wp.pdf.

Delay-Sensitive Cloud Computing and Edge Computing For Road-Safety Systems

Jan van de Kerkhof
Aalto University School of Science
jan.vandekerkhof@aalto.fi

Abstract

With the increase in vehicular communication technologies and the development of cloud and cloud-edge computing technologies, the capabilities of Intelligent Transportation Systems (ITS) increase and with them the possibilities of developing road-safety systems. This paper provides a comparison of two deployments of such systems, the first one being a vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) implementation that makes use of road-side units (RSUs) and the second one being an implementation that uses the LTE network, with cell caching at the base station, a technique that caches data that is in high demand. Both implementations make use of a (distant) cloud server. The comparison shows that the RSU implementation is more suitable with regards to capacity, and that the LTE implementation provides more potential for an easy deployment of the application. However, both implementations fail to provide an optimal implementation in both feasibility and costs with respect to the currently available technologies.

1 Introduction

During the last two decades, Intelligent Transportation Systems (ITS) have been developed to improve the quality and safety of commuting. Even though these systems have been contributing to a better flow of traffic and commuting experience, much room for improvement remains. There are still many traffic accidents around the globe and three fourths of these can be attributed to human error [12]. Moreover, studies have shown that 50-60% of traffic congestion in American cities is a result of traffic accidents [18]. Furthermore, there have been considerable developments lately in the area of self-driving vehicles. Google has unveiled its first build of their self-driving car prototype, for which the state of California has approved testing on public roads [2]. Also, a world-premiere testing of a self-driving bus without a steering wheel has been scheduled for December 2015 in Wageningen, the Netherlands [4]. All these examples explain the need for intelligent road-safety systems and applications. Some examples of these types of applications are, for instance, an emergency warning system, a lane change assistant or an intersection coordinator. Furthermore, these systems are able to detect traffic sign violations and give warnings about current road conditions. Recent developments in communication and computing technologies, such as the deployment of LTE and the shift in paradigm towards cloud-computing, enable development of these advanced applications. An example of an advanced ITS system is the ITS corridor that is being deployed between Rotterdam and Vienna, a first of its kind highway that is being equipped to fully autonomously shepherd cars without human intervention, coordinating traffic flow and warning about road constructions [1]. In order to implement an intelligent road-safety system, the infrastructure has to meet certain requirements. Firstly, since the system has to function in real-time in critical road-safety situations, safety messages need to be delivered and delay needs to be kept to a minimum as the latency deadline is strict. Secondly, there is the issue of capacity, as the system needs to be able to handle subscriptions to the size of entire cities.
The outline of this paper is as follows. Firstly, section 2 gives an overview of the requirements of road-safety systems. Then, sections 3 and 3.1 provide an overview of an implementation that uses RSUs, after which sections 3.2 and 3.2.1 provide an overview of the cloud and cloud-edge computing technologies that can be used for this implementation. Then in section 4, an overview is given of the alternative, LTE based, implementation, where section 4.1 provides an overview of the Nokia RACS technology that can be used for this implementation. After this, section 5 draws conclusions about the best implementation with the current technologies and makes speculations regarding possible future implementations. Finally, section 6 gives some suggestions for future work.

2 Requirements of road-safety systems

The main function of a road-safety system is to keep the drivers safe and avoid collisions. Keeping a driver safe can be done with different types of information and safety messages, some of which are more time-critical then others. Therefore, in road-safety systems there is a distinction between safety-critical messages and messages warning about the environment or road conditions, where a safety-critical message signals immediate danger. As stated by T.K. Mak et al. [11], S. Kato et al.[8] and Q. Mu et al.[17], the maximal delay for a safety-critical message in road-safety applications is 100ms. This is based on the assumption that each participant in the system broadcasts its positioning data at a frequency of 10Hz. This means that if a message is delayed by more than 100ms, the sender has already produced a new positioning message, rendering the previous one obso-

lete and inaccurate. In both the implementations discussed in this paper, safety-critical communications will have to flow through a centralized unit, the data center. The reasons for this are explained in detail in sections 3 and 4. Thus, for a road-safety application to function desirably, the round-trip-time (RTT) between the user equipment (UE) and the data center, plus the processing time of the application, will have to be less than 100ms.

3 Deployment with road-side units

One way of constructing the infrastructure required for a road-safety application is through Vehicular Ad Hoc Networks (VANETs), communication networks with fast-moving nodes (verhicles). They consist of on board units (OBUs) built into vehicles and road-side units (RSU) deployed on roads and highways that facilitate vehicle-to-vehicle communication (V2V) as well as vehicle-to-infrastructure (V2I) communication. V2V communication can be used for vehicles to broadcast traffic-related information (through periodic beaconing), such as vehicle location and speed, to all nearby vehicles. This raises better awareness of the surrounding traffic conditions. V2I communication can be used for traffic management purposes (e.g. lane-changing assistance) and to make participants aware of road-conditions (e.g. sharp turn ahead or a road blocked by an accident) [6]. The dedicated short range communication (DSRC) standard has been developed for the purpose of V2V and V2I communication [17] in VANETs. Furthermore, the WAVE (wireless access in vehicular environments) or IEEE 802.11p standard has been developed as an addition to the Wi-Fi standard to enable wireless communication amongst high-speed vehicles and between vehicles and road-side equipment. Even though vehicles can communicate with each other through DSRC, applications that solely rely on V2V communication face two major problems. The first is that it is hard to find a communications coordinator in a highly dynamic VANET, making the coordination of complex traffic situations in the fraction of a second highly challenging. The second of these problems is that the WAVE channel is memory-less, making it vulnerable to obstacles. This paper therefore assumes that any road-safety system requires both V2V communication and V2I communication (with an underlying data server as coordinator) for the system to function properly. As a side note, numerous issues still remain regarding VANETs that are currently being researched. These issues mainly regard the safety of the network against intruders and effective routing in such rapidly changing topologies, but this is beyond the scope of this paper. Figure 3 gives an example of a functioning VANET. One road-safety application that has been developed based on VANETs is SAFESPOT [5], which is an application that enhances a driver's awareness of the traffic situation by displaying appropriate warning messages. The system collects data from various resources, including other drivers and road-side infrastructure, and uses this to build a dynamic map of the current traffic situation. P. Jaworski et al. [7] propose a concept for an urban traffic management system that is based on an infrastructure that uses RSUs, the main goal of which is to maximize traffic flow and improve the overall safety of

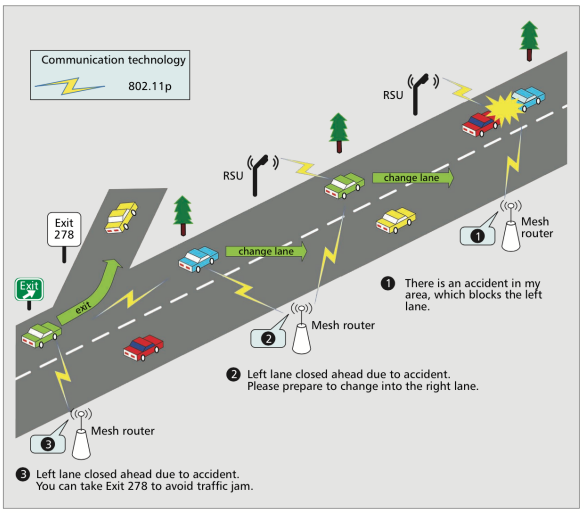

Figure 1: Example of a VANET [10]

traffic by managing traffic at every intersection in dense urban areas. The concept bases some of its functionality, such as data acquisition and the building of a dynamic map, on the SAFESPOT project. The authors design their system as a cloud/grid based computing system by dividing their application into six distinct interacting service layers, among which are an intersection control service and a routing service. This design shows that it is possible to develop a highly scalable cloud-like road-safety application based on the RSU infrastructure. The application in [7] is designed only for traffic coordination and not for safety-critical situations like collision prevention, but could be adapted to meet these requirements, as building such an application does not require any additional infrastructure.

3.1 ITS corridor

A real-world example of a system that relies on V2V and V2I communication is the cooperative ITS corridor [14] [1] that is being developed by The Netherlands, Germany and Austria. The work on it will start in 2016. It is a highway that is being equipped with ITS equipment and its purpose is to coordinate traffic by providing cars with speed recommendations, thus preventing the formation of traffic jams generated by the 'shock wave' effect of suddenly braking cars, by slowing down cars further down the highway. The system relies on the 802.11p (WAVE) Wi-Fi standard for V2V and V2I communication. It is possible that an LTE connection is going to be used to acquire probe data (location, speed) about every car, which will be used for traffic management. The ITS corridor does not provide real-time traffic-safety functionality. The testing of the system will be performed on a stretch of highway that is equipped with camera poles and Wi-Fi access points every 100 meters and Wi-Fi antennas every 500 meters. Since the WAVE standard only provides reliable communication of around 300 meters, the access points will have to be placed close to each other. The information that is gathered by the system is processed through central ITS stations. These stations are all responsible for a large

section of the road-network. Regarding the cloud-computing aspect of ITS systems, this central ITS station can be considered as the data center that will host the application. Thus, the ITS corridor, although it does not provide road-safety functionality, is almost a one-to-one fit with the theoretical cloud-computing models provided earlier, as it uses almost all of the same technologies and infrastructures.

3.2 Cloud computing

In order to provide the proper storage and computing capacity of a widely deployed road-safety system, a scalable, cloud-based computing platform is required, and there are several technologies and cloud-providers that will provide solutions. First, there is conventional cloud computing, where applications are run in a data center using highly flexible and scalable storage and processing resources. Here, it is also possible to make use of frameworks such as Google MapReduce, Apache Hadoop and Apache Spark to compute large amounts of data in parallel. However, these frameworks are more for running complicated queries than for communicating real time data. Therefore, analysis of these frameworks is beyond the scope of this paper. Secondly, there is the upcoming paradigm of cloud-edge computing [15], where virtual machines are offloaded to be close to the client in order to reduce latency between the UE and the data-center. Within this paradigm also falls the technology to place a RACS server [3] on the ENodeB itself, enabling cloud-computing over LTE with a highly reduced latency, which will be discussed in more detail in section 4.1. As established in the previous section, the maximum latency for criticial messages in a road-safety application is at most 100ms, and thus this is a good metric to analyse the feasibility of these cloud solutions and cloud providers with in terms of maximal latency.

3.2.1 Conventional Cloud Computing

When hosting an application in the cloud, we can choose a different number of cloud providers. Each of these providers has different qualities, like better computational or storage performance. In order to distinguish between these providers, A. Li et al. have developed *Cloudcmp* [9], a systematic comparator of the performance and cost of cloud providers. They compare data centers of the four biggest cloud providers, namely Amazon AWS, Microsoft Azure, Google AppEngine and Rackspace Cloudservers. One of the tests they run is a *ping* and *iperf* test to measure the optimal RTT between 260 global vantage points in the PlanetLab network. Their results, illustrated in figure 2, show that for at least one out of these providers, the RTT will always remain below 100ms, regardless of the vantage point positions. The more interesting results, though, are for the least performing cloud provider. As this cloud provider only has data centers in North America, there is a clear difference between the RTT of vantage points in North America and those from other continents. This turning point occurs around a RTT of 50ms, so we can therefore assume that for any data center that is in the same continent, the RTT should not exceed 50ms. This seems to be promising in the road-safety context, as latency can be kept under the threshold of 100ms.

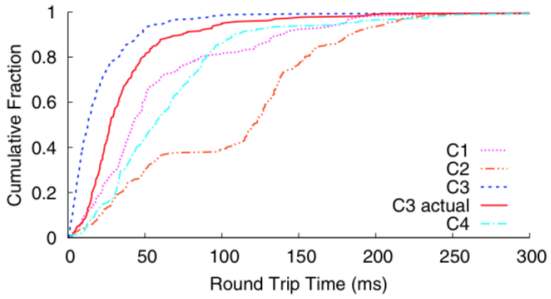

Figure 2: Results from *Cloudcmp* [9], where the RTT latency is measured from 260 different global vantage points

However, part of this functionality depends on the way the broadcasting of safety-critical information is implemented. If the broadcasting can be done by not (persistently) storing anything and working out of cached data, then this performance is sufficient for the road-safety application, as the processing time in the data center will be negligible. However, if the broadcasting of locations or any other safety critical function of the application require access to a persistent data store, then some cloud providers might not provide the right level of latency performance. In *Cloudcmp*, tests are also run to test the average response time of the database services provided by these providers, and they show that for a `get` operation on a table with around 100K entries, the response time is over 100ms for the 95th percentile of all `get` operations. Thus, for any solution that would require persistent storage of a large number of records that need to be accessed for real-time safety messages, *all* of the cloud services would not provide the required response time. One of the benefits of cloud computing is the scalability of the application. Both application discussed in this paper rely on a cloud application for computation and therefore do not have to worry about computational capacity requirements, as the capacity of any cloud application is easily scaled.

3.2.2 Cloud-Edge Computing

Lately there has been a paradigm switch towards cloud-edge computing in order to enable low-latency access to the cloud, as WAN latencies are not likely to improve in the near future. Cloud-edge computing is the process of exporting an VM-based image of the cloud to a local *cloudlet*. A cloudlet is a resource-rich machine or multiple machines close to the client, which only caches a copy of the cloud-based applications' code. It is basically a "datacenter in a box", which allows for low-latency, one-hop access between the user and the cloudlet based on (wireless) LAN access [15]. This virtual machine needs to be downloaded and installed to the cloudlet, which is achieved between 60 and 90 seconds on average, after which the client is able to connect to it. This technology could be used to reduce the latency between the RSUs and the data center, by constructing nearby cloudlets in any area that implements the system, where every cloudlet runs a VM-based representation of the overall system. This way, the latency between the cloud and the RSU can be reduced to almost nothing, as the cloudlet is located in the di-

rect vicinity of the RSUs. These cloudlets do however add another component to the system that runs up the overall deployment cost of the system.

4 LTE as an enabler of road-safety applications

Apart from the previously researched implementation that utilizes road-side units, LTE also provides much potential as an enabler of road-safety systems. Instead of relying on V2V and V2I communication, all of the communication can be done through LTE, cell towers and a centralized processing unit, the data center. Deploying a road-safety system through LTE is therefore relatively easy, since the infrastructure already largely exists and is deployed worldwide, thus there is no need to deploy any (costly) road-side infrastructure. Additionally, LTE could be used not only autonomously but in combination with V2I solutions, for instance to provide V2I communications in areas where RSUs are not densely deployed. However, LTE has two major challenges that come with it. The first of these is *latency*. Since LTE does not support ad-hoc communication between UEs, all the communication in the system has to go through the cloud server that is connected to the user through the ENodeB and the Evolved Packet Core (EPC). The latency of a message is thus the product of the round-trip-time (RTT) between the UE and the ENodeB and the RTT between the ENodeB and the server. The second limitation is *capacity*. Cellular networks are typically not optimized for a large number of users sending frequent, small, latency sensitive messages.

In [8], S. Kato et al. explore these two main limitations in the context of road-safety applications. In order to classify the system by means of Quality of Service (QoS) (which in essence means latency), they introduce the concept of *data freshness*, an upper boundary on message delay that is required for the system to function as desired. As mentioned before, they set this data freshness to 100ms. The authors point out that in a road-safety scenario, the network utilization grows exponentially as the number of UEs increases. This is due to the fact that for an unconnected device in an LTE network, it takes a minimum of 100ms for the ENodeB to allocate resources to the device. Therefore, in order to achieve a data freshness of 100ms, all the devices in the application need to be constantly connected to the network, which highly increases the network utilization.

The authors evaluate the performance of LTE in terms of the maximal number of UEs that are supported and the utilization that the road-safety application has in terms of the total bandwidth. They make a distinction between MBMS and non-MBMS enabled LTE. In the MBMS (Multimedia Broadcast/Multicast Service) enabled mode, it is possible to broadcast messages to several UEs. This means that location information on nearby vehicles, which normally has to be transmitted to every UE individually, can be broadcasted as a single message. This reduces the number of messages that are transmitted across the network by one or two orders of magnitude. Unfortunately, this feature is not enabled in many base stations and would therefore not be a feasible option in many areas, without adjusting the settings at the base station. Finally, the authors look at the effect the location of the data center has on these parameters. Since placing the data center close to the ENodeB highly reduces latency, the network utilization drops significantly. In the optimal case, where the data center is located at ENodeB, the theoretical maximal number of UEs that can be supported in non-MBMS mode is 73 with a data freshness of 100ms versus 162 at a 200ms data freshness. The utilization of the network is very undesirable in both cases, with a 77.5% and 86.4% downlink utilization, respectively. In the MBMS case, the results are much more promising, where the maximal number of UEs is 3076 with a data freshness of 100ms and 6783 at a data freshness of 200ms. The utilization is 14.4% and 4.6%, respectively.

Now that the theoretical limit on the number of supported UEs is known, it makes sense to look at the average number of vehicles that would have to be supported per ENodeB, as this directly influences the feasibility of the deployment, since the system needs to be able to support all the vehicles in the area. In [13], T. Mangel et al. research the cell density for the city of Munich with respect to the amount of traffic in the city. A cell consists of the ENodeB together with the corresponding antennas. They conclude that "the cell size in Munich is roughly $0.41km^2$ [...] the average expected amount of vehicles per cell is $\approx$50. However, maximum values turn out five to six times higher. Rush hour values will be even higher [...]". Thus, it makes sense to assume a vehicle density of around 300 for busy areas in a city and to assume this value much higher during rush hours, according to the authors this can be around 600 vehicles. If we look at the maximum supported number of vehicles in non-MBMS mode, which is 73, this value is exceeded greatly, and the system could only be supported in MBMS mode, where the theoretical maximum is 3076. However, this requires all of the base stations to enable MBMS mode. A final thing that has to be taken into account is that the network utilization of this solution will be significant and will have to be paid for, which will result in a high cost to the users.

The limitations of LTE in terms of communications capacity and of network cost do not exist for the implementation with RSUs, where the communications capacity is sufficient for any traffic situation and the network usage does not have to be paid for by the end users.

4.1 Nokia RACS for LTE

Related to the previously mentioned advancements in cloud-edge computing, there have also been developments in the area of cell caching [16], the technology that will enable the base stations of each mobile cell to cache popular content. This is motivated by the change in network usage, where many users download articles and videos through the wireless link. This vast increase in traffic increases the load on the infrastructure behind the base stations. One of the solutions to this problem that is currently being developed is the Radio Access Cloud Server (RACS), that is being developed by Nokia and Intel. The RACS server allows for the development of so called "Liquid Applications", VM-based applications that run at the base station and that cache popular content. While this development has been mainly motivated

by the massive growth of video traffic [3], this technology could be used for the development of a road-safety application, where instead of popular video content, the location of all the vehicles in the area and current road-conditions can be cached. By caching this data at the base station, the only latency bottleneck that the application will experience is the radio interface itself. By using the RACS server or any other form of cell caching that would provide the same functionality, we could achieve the theoretical optimal limit described by S. Kato et al. [8] in terms of capacity, as the data center, or at least the data center functionality, would be located right at the ENodeB.

5 Conclusions

5.1 Best fit for the current economic climate

This paper has looked at two different implementation of road-safety systems considering the current available technologies. Both implementations have their upsides and their downsides. The first implementation, employing V2V communication, V2I communication and RSUs, will meet the requirements in terms of latency and capacity, when deployed in the cloud through a datacenter that is located on the same continent. If a very strict maximal latency is required, the system could also make use of cloud-edge computing to include nearby cloudlets that reduce latency to a minimum. However, this implementation is costly with regard to the deployment of RSUs and possibly the cloudlets, but that the system is feasible is being shown by the deployment of the ITS corridor. The alternative, a deployment over the LTE network combined with cell caching, is feasible in terms of latency, but will suffer from problems with regard to capacity, as the amount of supported UEs is insufficient in non-MBMS mode. However, in MBMS mode, the capacity should be sufficient, although this is not enabled in many base stations. This solution, also, will be costly in terms of data usage across the LTE network. Thus, both of the options are theoretically feasible, but are far from optimal in the current economic climate.

5.2 Future possibilities

Looking to the future of cellular networks, one of the most limiting factors right now is the capacity of the LTE network and the time it takes to allocate resources to an idle UE. If the capacity of the cellular network increases and resource allocation becomes faster, for instance by 5G, the feasibility of a solution through the cellular network becomes much higher as much more UEs would be able to connect to the network. Also, the load on the network would decrease and in turn the cost of using the network would decrease. Therefore, it is assumable that as the capacity of cellular networks increases, the possibilities of implementing a road-safety system at low cost become much larger. Looking at the future of RSU deployment is not sensible, as this is already a realistic possibility.

Furthermore, this paper make the assumption that intelligent vehicles in the future would require an extensively deployed road-safety system to function, whereas the development of the self-driving car give thought to the possibility of autonomously functioning vehicles that would use cameras and sensors to drive, and that would not depend on a road-safety system to function.

6 Future work

In this paper the assumption has been made that a road-safety needs a centralized coordinator and that such a system cannot function on V2V communications alone. Future work could be to develop a road-safety system that would function solely on DSRC, using the WAVE standard. Also, one could research the long-term cost difference between a deployment with RSUs and a deployment over the LTE network.

References

[1] Cooperative ITS Corridor Deployment, url = `http://www.bmvi.de/shareddocs/en/anlagen/verkehrundmobilitaet/strasse/cooperative-its-corridor.pdf?__blob=publicationfile`, note = Accessed: 2015-02-02.

[2] Google unveils 'first real build' of its self-driving car prototype, url = `http://mashable.com/2014/12/23/google-first-real-build/`, note = Accessed: 2015-02-02.

[3] Increasing Mobile Operators' Value Proposition With Edge Computing, url = `https://networkbuilders.intel.com/docs/nokia_solutions_and_networks.pdf`, note = Accessed: 2015-03-19.

[4] Selfdriving car on road in December, url = `http://nos.nl/artikel/2015674-zelfrijdende-auto-in-december-de-weg-op.html`, note = Accessed: 2015-02-02.

[5] F. Bonnefoi, F. Bellotti, T. Scendzielorz, and F. Visintainer. Safespot applications for infrasructurebased cooperative road safety. In *14th World Congress and Exhibition on Intelligent Transport Systems and Services*, pages 1–8, 2007.

[6] H. T. Cheng, H. Shan, and W. Zhuang. Infotainment and road safety service support in vehicular networking: From a communication perspective. *Mechanical Systems and Signal Processing*, 25(6):2020–2038, 2011.

[7] P. Jaworski, T. Edwards, J. Moore, and K. Burnham. Cloud computing concept for intelligent transportation systems. In *Intelligent Transportation Systems (ITSC), 2011 14th International IEEE Conference on*, pages 391–936. IEEE, 2011.

[8] S. Kato, M. Hiltunen, K. Joshi, and R. Schlichting. Enabling vehicular safety applications over lte networks. In *Connected Vehicles and Expo (ICCVE), 2013 International Conference on*, pages 747–752. IEEE, 2013.

[9] A. Li, X. Yang, S. Kandula, and M. Zhang. Cloudcmp: comparing public cloud providers. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, pages 1–14. ACM, 2010.

[10] X. Lin, R. Lu, C. Zhang, H. Zhu, P.-H. Ho, and X. Shen. Security in vehicular ad hoc networks. *Communications Magazine, IEEE*, 46(4):88–95, 2008.

[11] T. K. Mak, K. P. Laberteaux, and R. Sengupta. A multi-channel vanet providing concurrent safety and commercial services. In *Proceedings of the 2nd ACM international workshop on Vehicular ad hoc networks*, pages 1–9. ACM, 2005.

[12] L. Malta, C. Miyajima, and K. Takeda. A study of driver behavior under potential threats in vehicle traffic. *Intelligent Transportation Systems, IEEE Transactions on*, 10(2):201–210, 2009.

[13] T. Mangel and H. Hartenstein. An analysis of data traffic in cellular networks caused by inter-vehicle communication at intersections. In *Intelligent Vehicles Symposium (IV), 2011 IEEE*, pages 473–478. IEEE, 2011.

[14] P. Ross. Thus spoke the autobahn. *Spectrum, IEEE*, 52(1):52–55, 2015.

[15] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies. The case for vm-based cloudlets in mobile computing. *Pervasive Computing, IEEE*, 8(4):14–23, 2009.

[16] X. Wang, X. Li, V. C. Leung, and P. Nasiopoulos. A framework of cooperative cell caching for the future mobile networks.

[17] Q. Xu, T. Mak, J. Ko, and R. Sengupta. Vehicle-to-vehicle safety messaging in dsrc. In *Proceedings of the 1st ACM international workshop on Vehicular ad hoc networks*, pages 19–28. ACM, 2004.

[18] J. Zhang, F.-Y. Wang, K. Wang, W.-H. Lin, X. Xu, and C. Chen. Data-driven intelligent transportation systems: A survey. *Intelligent Transportation Systems, IEEE Transactions on*, 12(4):1624–1639, 2011.

More ICT to Make Households More Green

Hylke Visser
Student number: 467588
hylke.visser@aalto.fi

Abstract

The world has become increasingly concerned with the fast growing energy demand in households. A large portion of household energy consumption is wasted on devices that are unnecessarily switched on. Green energy sources are usually dependent on the weather, and require more flexible energy consumption. ICT can help sloving both of these issues.

This paper gives overviews of ICT-based methods that address these issues. It describes the concept of home automation and shows how this can help save energy, while also becoming more flexible in energy consumption. This paper also describes how ICT can help change energy consumption behavior in households. Finally, the authors propose a direction for future research.

KEYWORDS: ICT, household energy consumption, energy saving, home automation, behavioral change

1 Introduction

Global energy demand is growing extremely fast. Increased urbanization, industrialization and the growing economies in developing countries are expected increase energy demand even more. In recent years, the world has become increasingly concerned about the consequences of our energy consumption. Ensuring that this fast growing energy demand can be satisfied by the current energy production is a challenging task.

Climate change has raised the demand for more sustainable energy sources such as wind power, hydro-power and solar power. Unfortunately the energy production capacity of these sources is highly dependent on the weather and other environmental influences. This requires more flexible energy consumption to avoid the need for back-up systems. These back-up generators supply extra energy during peak demand and are often fossil fuel based systems. ICT will play a big role in helping energy consumers be more flexible. It will help mitigating these peaks and reduce the need for the back-up systems.

A surprisingly large percentage of energy consumption is attributable to households. In the United States, this percentage is 22% [1] and in Europe 26% [2]. In China, the industry sector currently the largest consumer of energy, and households account for only 10%[14]. This is however expected to change very fast in the next years as the latter will raise their demand. The U.S. Department of Energy[3] predicts that in the near future, households will be responsible for almost 40% of the annual energy consumption. Households are therefore a good target for reducing energy consumption and increasing the flexibility of this consumption.

A large part of energy is lost between the energy source and home appliances. This means that reducing unnecessary consumption of energy can have an even bigger impact than one would think. Clever ICT solutions can give consumers insight into their energy consumption and motivate them to waste less energy.

This paper is structured as follows. Section 2 describes the layers and components of a typical home automation environment in a bottom-up approach. It describes the entire process from collecting measurements to making decisions and gives for each step in this process an overview of the different solutions that are currently used or have been proposed. Section 3 gives an overview of different strategies for changing human energy consumption behavior. It also describes the ways in which ICT can help giving household owners insight in their energy consumption and change their energy consumption habits. Section 4 lists a number of practical solutions that have been deployed commercially and discusses the relation between the solutions and the previously discussed scientific research. Section 5 argues about the difference between home automation solutions and behavioral change applications. It shows why these solutions could work or why they fail. Section 6 gives the conclusions of this paper and proposes a future direction for research.

2 Home Automation

This section gives an overview of the methods and technologies that are used to facilitate energy saving in a household using home automation.

A typical home automation system is very similar to the architecture presented in 2003 by Cook et al. [7]. This architecture consists of several layers, each having their own functions and responsibilities. The physical layer contains hardware that measures energy consumption. The communication layer facilitates data exchange between nodes in the sensor network and devices that collect and process this data. The information layer gathers, stores and aggregates sensor data into knowledge bases that can be used by smart decision-making systems. The decision layer determines the actions that have to be executed to transition the environment into the desired state.

2.1 The Physical Layer

Measuring energy consumption requires the installation of sophisticated sensors. Mattern et al. [10] describe two dis-

tinct approaches for the installation of sensors. Each of these has its advantages and disadvantages.

A single sensor approach is a low-cost solution that usually requires low initial effort. The simplest solution requires nothing more than installing a single sensor on the existing energy meter. The downside of this approach is that it does not provide detailed information as its accuracy is limited to the energy consumption of the total household. It is therefore not easily possible to give consumers concrete recommendations on which devices he can switch off to save energy. By using sophisticated pattern recognition systems as described in Section 2.3, it is possible to make assumptions regarding the cause of fluctuations in the total energy consumption of a household. This could be used to provide more details about the energy consumption of specific devices.

A multiple sensor approach can be used to monitor the energy consumption of individual appliances. This makes it easier for the consumer to see which devices consume energy and, in turn, make decisions based on this information. Unfortunately, this approach is expensive, as it requires many energy sensors. The approach could also decrease the accuracy at a household level, as some appliances might not be attached to an energy meter.

One solution for this would be to install sensors at a higher level, for example in every room of a house. Together with contextual sensors that measure temperature, light and presence of people, this approach would provide valuable data for home automation.

2.2 The Communication Layer

Using an approach that uses multiple sensors requires combination and aggregation of the data that is gathered by those sensors. In many studies a smart gateway is responsible for this task. The first step is to communicate energy measurements from the sensors to these gateways.

Kovatsch et al. [9] give a good overview of the different standards for communication in home automation environments. The X10 system[1] has existed since the 1970s and uses AC power lines for communication. KNX[2] is a newer system and uses twisted pair cables for communication. DigitalSTROM[3] also uses the existing 230V mains.

Recent studies seem to be more confident about wireless standards such as ZigBee and IPv6. ZigBee is a IEEE 802.15.4 based standard that provides inexpensive, low-power, wireless communication. ZigBee is one of the most used systems for Wireless Sensor Networks, and also very suitable for Home Automation.

Weiss et al.[12] used smart power outlets, called Ploggs[4] to build a smart home environment. These Ploggs can communicate over Bluetooth and ZigBee, which allow for low-power, short range communications of energy consumption data.

Kovatsch et al. show that IPv6 is a suitable alternative to the previously discussed standards. The Internet Protocol Suite has proven to be a matured networking concept. The rapid growth of the Internet is proof that IPv6 is extremely scalable, which is exactly what is needed for Home Automation and the Internet of Things. The 6LoWPAN[5] standard allows resource-constraint devices to run the IPv6 stack. Installation of IPv6-capable devices is easy and can benefit from the autoconfiguration mechanism of IPv6. Also, IPv6 enables smart appliances to integrate with the Web, which allows for natural user interaction. Also, the Internet already provides many well-established security mechanisms that can be used to secure these Home Automation systems, but Kovatsch et al. also mention that more research has to be conducted in the area of security.

2.3 The Information Layer

The data that is collected by the network of sensors can be collected and aggregated by gateways. These devices will play a central role in home automation. Gateways will analyze the measurements made by the sensor network that has been deployed in the house. In addition to the previously discussed sensors for measuring energy consumption, this network also contains other sensors. Temperature sensors are used to make decisions on whether to activate or deactivate the heating or cooling system for a room. Light sensors can help determine how much brightness the lights in the room should add to the natural daylight. Motion sensors can help avoid unnecessary energy consumption by turning off devices in empty rooms. Many other metrics can and will be added in the future, and the gateway will have to interface with all of these.

By analyzing the measurements of the sensor network over a longer time span, the gateway can build an accurate model of the household. Building knowledge about individual appliances can be done with sophisticated pattern recognition or deep learning algorithms. The eMeter system by Mattern et al. [10] makes this process as simple as pressing a "record" button and switching the appliance on and off. The system will detect the change in current and store this to its database.

Probably the most important knowledge for home automation is the presence of people. If the home automation system knows about the habits and schedules of the inhabitants of a house, it can make energy saving decisions. Barbato et al. [5] proposed a system that creates different profiles based on sensor data. For the user presence profile, the system aggregates 24 hours of data for each room in the house to build a daily profile in a given monitoring period (a week). Similar daily profiles are clustered together.

At any time the user can manually regulate the light and temperature in a room. This information is then recorded by the system to create a profile of the users' preferences.

2.4 The Decision Layer

The decision layer is responsible for activating or deactivating devices in the smart home. By combining the user's preference profiles with the profiles derived from sensor data, the decision layer can predict and execute the necessary actions needed to satisfy the user's preferences. This could

[1] x10.com
[2] knx.org
[3] digitalstrom.com
[4] Website offline. Old website can by found by searching for ploggintern ational.com/ploggproducts.html on web.archive.org
[5] ietf.org/rfc/rfc4944

also allow the system to save energy by opening curtains in the morning, which would result in the user not having to switch on the lights, or by activating the heating system before the user comes home, so that the user does not have to set the heating system to maximum on arrival (which would also require relatively more energy).

These last examples demonstrate well the possibilities for balancing the load on the power grid. This load balancing will be essential when electricity is generated by green sources, as the supply of most of these will be highly dependent on the weather. Smart homes can adapt to this dynamic energy supply and decide when to switch on washing machines, lower the temperature freezers or when to charge electric cars. This system also benefits consumers who have dynamic energy pricing, which will be used in the future smart grid.

3 Behavioral Change

Home automation is a step in the right direction, but even more energy can be saved by addressing the human factors in energy consumption. People can have many reasons for saving energy. Saving the planet, a feeling of moral obligation or the desire to save money may lead to greater attention to someone's energy consumption behavior. Unfortunately, many households hardly pay attention to their energy consumption. However, there are various methods aimed at changing the energy consumption behavior of households.

Abrahamse et al. [4] give a good overview of different strategies for informing households and changing their energy consumption behavior. They divided the different methods into two categories: antecedent interventions and consequence interventions.

The most simple form of an antecedent intervention is information. Mass media campaigns aimed to change consumers' behavior seem to have almost no influence. Workshops did change the level of concern of the attendees, but did not actually change their behavior. The best information-based intervention method is giving consumers tailored information through home audits. More successful antecedent interventions included some form of goal setting: household owners or external parties can set a goal for saving a certain percentage of energy. This goal can even be converted to an oral or written pledge or promise to save energy, which is then called a commitment.

Consequence interventions are based on the idea is that positive consequences will encourage desirable energy consumption behavior. Direct and continuous feedback about household energy consumption can influence behavior. Households can immediately relate their actions with the outcomes. This continuous feedback can be given using a display that shows the current energy consumption or the costs of the current energy consumption. Comparative feedback can be even more effective. Seeing your performance relative to the performance of similar households can lead to a feeling of competition or social pressure. Energy saving then becomes a sort of a game, motivating people to save even more energy. An extra motivator for people to save energy are rewards. Monetary rewards can be either based on the amount of energy saved or can be a fixed amount when a certain target has been reached.

3.1 Energy Consumption Feedback

The most important factor in changing the behavior of energy consumers is the feedback on their energy consumption. Many different methods exist for providing this feedback. Karjalainen et al. [8] have analyzed different ways of presenting feedback on energy consumption.

The most simple form of feedback they mention is simply presenting the energy consumption of the household. This information can be presented as the number of kWh consumed in a period of time. More relevant information can be the cost per hour or the environmental impact. Disaggregating energy consumption by times of the day, by room or by appliance significantly improves the relevancy of the information and helps users identify appliances and habits that consume a lot of energy.

Comparisons can give even better insights. When someone uses goal-setting to reduce their energy consumption, it is important for them to compare their current energy consumption to the goal. By storing the energy consumption data on a hard disk, it is possible to build historical comparisons. However, Karjalainen et al. stress the need for normalization of these comparisons, as the weather can have a large influence on energy consumption.

A normative approach compares the energy consumption of a household to that of other households. This can be performed on a national or regional level, or even within the same neighborhood. The relevance of the comparison highly depends on the similarity between the households that are compared. This method can motivate people to save more energy, however, Karjalainen et al. note that people might save less if they see that they "have saved enough" or will justify their higher consumption with excuses instead of trying to change it.

In a study by Bonino et al. [6], many of the previously mentioned methods are combined. Their experiment showed a display with the current power consumption, the total power consumption for that day and the goal the user had set previously. Furthermore they showed a map of the house, with each room colored green, orange or red, depending on the energy consumption in that room. A large number of people participated in a survey that confirmed that such feedback is indeed understandable and motivates behavioral change.

Energy consumption feedback is also used to show how much energy is consumed by different appliances in the household. This information can then be used to determine whether replacing appliances can save energy (and money). The PowerPedia system, created by Weiss et al. [13], helps users better understand the power consumption of their appliances. After measuring the consumption of a device, the results are added to a global database. The user can compare his results to similar devices that have been measured by other people using the PowerPedia system.

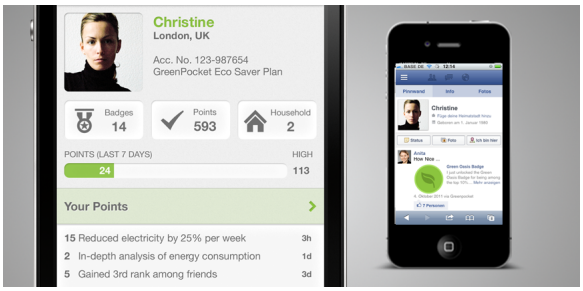

Figure 1: GreenPocket's social metering application
source: greenpocket.de/en/products/residential-customers/social-metering

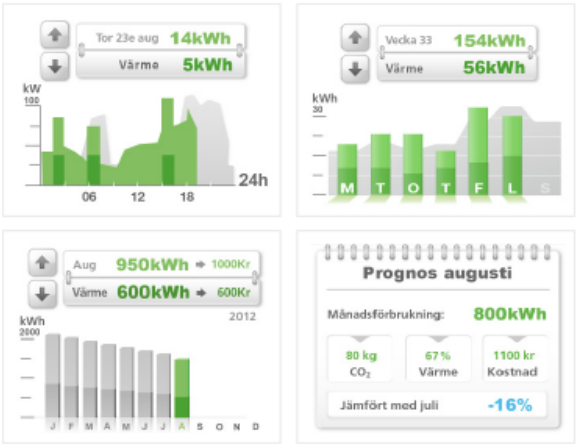

Figure 2: ELIQ's Energy Display
source: eliq.se/en/products/eliq-energy-display

4 Practical Solutions

GreenPocket[6] is a German company that enables utilities to leverage the power of smart meters and smart homes. Their Energy Expert Engine cobmines energy consumption data with external data such as the weather and then builds personalized data based on customer data and information about individual households. A visualization of this data is accessible by the household owners on a web interface and in mobile applications.

GreenPocket also has a mobile application (see Figure 1) that adds a social component to energy consumption. With this application energy consumers can participate in a competition that motivates them to reduce their energy consumption.

The GreenPocket Smart Home solution enables users to control their smart home by using pre-defined schedules and rules, but also allows them to use their smartphones to manually control their home. Together with its other solutions, GreenPocket provides a powerful system that encourages energy consumers to change their behavior and even helps them to do that.

Another company, Opower[7] conducted a multi-year study into the relationship between energy consumers and utilities. From this research they defined the *five universal truths*[8] about utilities and energy consumers:

- "Utilities are not meeting customer expectations"
- "Everyone wants lower bills"
- "People look to utilities for energy information"
- "Customers value personalized energy insights"
- "Everyone wants to know how they measure up"

Based on these "five universal truths", Opower provides a service that helps utilities better satisfy their customers. They give consumers personalized and detailed energy insights, together with tips about how they can lower their energy bill. Opower also lets consumers compare their energy consumption to that of similar households.

[6] greenpocket.de/en/
[7] opower.com
[8] opower.com/fivetruths/

On the website social.opower.com Opower provided an application where users could compare their energy consumption to the consumption of their Facebook friends. After eveluating three years of usage data, Opower decided that this application was not the most successful way to help customers save energy and they decided to take down the application.

ELIQ is a company that does not work with utilities, but provides a solution that allows household owners to install a smart metering system by themselves. The ELIQ system consists of an energy sensor that can be attached to the energy meter in the household (that is currently only used for billing). This sensor transmits the data to an in-house display or to an ELIQ Online device that enables the user to view all his data online.

Figure 2 shows a visualization of a household's current energy consumption together with historical energy consumption and predictions for the future. The system can also calculate the amount of energy that is consumed for heating the household to give better insights in environmental influences.

Selvefors et al.[11] evaluated ELIQ's system in 23 households located in and near the city of Gothenburg in Sweden. In a period of two years energy consumption data was gathered from both households that participated in the experiment and from a large sample of comparable households. The study shows that households that actively participated in the study and used ELIQ's system regularly achieved a significant reduction in energy consumption, while households that did not regularly use ELIQ's system did not see any reduction in energy consumption either. There are many possible reasons and explanations for the fact that many households did not use the application, even when it costs little effort and is easily accessible from the users' phones. Lack of time or interest are the main reasons for the apparent lack of success with similar applications, while the applications do have a large impact when they are actually used.

ELIQ also contains a social component that adds gamification to energy saving. By creating a network of "energy friends" it is possible to compare your energy consumption

with others. Games and challenges lead to a competitive way of energy saving.

5 Discussion

It is clear that there is a need for the systems and techniques discussed in this paper. Reducing unnecessary energy consumption and becoming more flexible with energy consumption will be extremely important in the near future. There is a large amount of studies that propose solutions for every step of the process. Energy sensors and smart sockets can detect unneccessary consumption. Many protocols and architectures are proposed for for communicating and aggregating sensor data. Smart algorithms can make decisions when this data is combined with knowledge about the users and facilitate sophisticated homa automation services. By visualizing energy consumption in a household and comparing that to a pre-defined goal or to consumption in other households it is possible to change user behavior.

Home automation systems are popular and many people are interested in those systems. In the future, Load Balancing and integration with the Smart Grid will be two very important features of a home automation system. Currently, however, the price and installation effort is high and slows down deployment of these systems.

Applications that are aimed at changing energy consumption behavior have a lower barrier, which is why these are more widely deployed. However, it looks like these systems do not meet their expectations. Most consumers are not genuinely interested in changing anything about their energy consumption. And people who actually start using energy saving applications, quickly lose interest after a few months.

Behavioral change is extremely difficult, especially when it concerns something that the majority of your target group is not really interested in. Forcing people to use an application that makes them do things they normally would not do, will give most of them an uncomfortable feeling.

Maybe the best way to save energy is not changing the user by showing him what he does wrong and should do instead. Maybe we should first focus on making the home itself smart, so that it can make the user feel more powerful, while saving energy without the user noticing. Home automation and energy behavioral change seem to be two different things, but in fact they complement each other. Home automation can be the first step towards changing the users' behavior. A smart home can detect possibly unneccessary energy consumption and ask the user what it should do. This gives the user control over his home, but simultaneously saves energy.

6 Conclusion

In this paper we presented a general overview of two methods for helping households save energy. Home automation introduces an infrastructure into the household that measures, collects and aggregates energy consumption data. Together with information from the environment, user presence profiles and user preferences this system can save energy by anticipating user actions and avoiding unneccessary energy consumption.

By showing detailed energy consumption information to a user, it is possible to give him more insight into his energy consumption. This can lead to a change in his behavior. By setting goals or creating challenges in the user's social group, it is possible to stimulate competitive energy saving behavior.

These methods work well, but adoption by consumers is still low. This is due to high cost and effort with home automation systems and relatively low interest for applications that stimulate behavioral change. Future research could therefore aim at integrating devices and appliances with smart home systems. This will lower the barrier to create smart home environments and create a powerful medium to reach users. With this medium it is possible to apply behavioral change methods in a way that feels less awkward than with current applications.

Nevertheless, saving energy in households and becoming more flexible in energy consumption is extremely important. Research into this topic will grow as energy consumption increases and with the expectations of the Smart Grid. And ICT will definitely play an important role in making households more flexible and more green.

References

[1] U.S. Energy Information Administration. Annual Energy Review 2011.

[2] European Environment Agency. Final energy consumption by sector.

[3] U.S. Department of Energy. 2008 buildings energy data book.

[4] W. Abrahamse, L. Steg, C. Vlek, and T. Rothengatter. A review of intervention studies aimed at household energy conservation. *Journal of environmental psychology*, 25(3):273–291, 2005.

[5] A. Barbato, L. Borsani, A. Capone, and S. Melzi. Home energy saving through a user profiling system based on wireless sensors. In *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 49–54. ACM, 2009.

[6] D. Bonino, F. Corno, and L. De Russis. Home energy consumption feedback: A user survey. *Energy and Buildings*, 47:383–393, 2012.

[7] D. J. Cook, M. Huber, K. Gopalratnam, and M. Youngblood. Learning to control a smart home environment. In *Innovative Appl. of Artifical Intelligence*, 2003.

[8] S. Karjalainen. Consumer preferences for feedback on household electricity consumption. *Energy and Buildings*, 43(2):458–467, 2011.

[9] M. Kovatsch, M. Weiss, and D. Guinard. Embedding internet technology for home automation. In *Emerging*

Technologies and Factory Automation (ETFA), 2010 IEEE Conference on, pages 1–8. IEEE, 2010.

[10] F. Mattern, T. Staake, and M. Weiss. Ict for green: how computers can help us to conserve energy. In *Proceedings of the 1st international conference on energy-efficient computing and networking*, pages 1–10. ACM, 2010.

[11] A. Selvefors, I. Karlsson, and U. Rahe. Use and adoption of interactive energy feedback systems. *Proceedings of IASDR*, pages 1771–1782, 2013.

[12] M. Weiss and D. Guinard. Increasing energy awareness through web-enabled power outlets. In *Proceedings of the 9th International Conference on Mobile and Ubiquitous Multimedia*, page 20. ACM, 2010.

[13] M. Weiss, T. Staake, F. Mattern, and E. Fleisch. Powerpedia: changing energy usage with the help of a community-based smartphone application. *Personal and Ubiquitous Computing*, 16(6):655–664, 2012.

[14] X. Zhao, N. Li, and C. Ma. Residential energy consumption in urban china. Technical report, 2011.

Software market of network functions virtualization

Aarno Vuori
Student number: 60104J
aarno.vuori@aalto.fi

Abstract

Mobile operators are going to need new ways to reduce their costs in telecommunications business. New technologies such as network functions virtualization (NFV) combined with cloud computing technology are going to be one solution for this in the future. The European Telecommunications Standards Institute (ETSI) and Group of Major Telecom Operators has formed the Industry Specification Group (ISG) for NFV. Today, this large community of experts is developing the required standards for (NFV). Mobile technology companies such as NSN and Ericsson are both involved in this group. The vision of this group is an open ecosystem for NFV which enables rapid service innovation for network operators and service providers. All services are enabled by software-based deployment by using virtualized network functions (VNF). The ISG maintains core NFV documentation, including an architectural framework and associated technical requirements. NFV documentation provides the technology platform to the NFV ecosystem. This enables open innovation in the design of VNFs and end-to-end networks. It will also reduce power usage for the network operator. This standardization is an ongoing project and it is going to drive the NFV work forward across the industry. Nokia Solutions and Networks (NSN) and Ericsson are both involved in NFV evolution. Both companies have developed products which use the NFV framework in their products. Ericsson's product is Ericsson Cloud System and NSN's product is VoLTE based on NFV architecture. The main purpose of this paper is to highlight above the mentioned new technologies and give examples of related new implementations.

KEYWORDS: core network, virtual evolved packet core (vEPC), commercial off-the-shelf equipment (COTS), infrastructure as a service (IaaS), hardware as a service (HaaS), platform as a service (PaaS), software as a service (SaaS), network as a service (NaaS), virtual network functions (VNFs), evolved packet core (EPC), Quality of Service (QoS), Quality of Experience (QoE), Long-Term Evolution (LTE) network.

1 Introduction

Increased demands for cheaper and faster network technology is a challenging task. Telecommunication companies struggle with the increasing cost of dedicated network hardware and declining revenues. Therefore, there is a clear need for new technologies to provide increased network performance and value with lower costs.[18] One proposed candidate for this is Network Functions Virtualization (NFV). NFV and Cloud Computing together enable cost effective mobile network functions. Future technologies such as 5G networks have a major effect on faster connectivity speeds. All this will need more financial investments. NFV is going to move the telecommunications core infrastructure towards more cost-efficient core.[8] A group of major telecom operators has formed an industry specification group for NFV under the European Telecommunications Standars Institute (ETSI). Recently, more telecom equipment providers and IT specialists have joined the group.[8]

NFV is a model and a technology which transfers network functions from dedicated hardware to software-based applications. These applications are consolidated and executed on standard IT platforms. Platforms consist of high-volume servers, switches and storages. Based on NFV, networks functions can be placed in various locations, for instance datacenters, network nodes, and end-user premises depending on what the network requires. [16] NFV can provide many advantages to the telecommunication industry: architecture openness of platforms, scalability and flexibility, operating performance improvement, shorter development cycles, and also reduced capital expenditure (CAPEX) and operational expenditure (OPEX) investments.[16]

Cloud computing is a quite new business model for operators and it gives a new way to offer service delivery. They can combine telecommunications and internet applications from cloud for services offered to consumers and enterprises. On this cloud computing model, telco companies must think privacy and security issues. These issues require careful business- and technology planning to meet all the needed requirements.[10]

2 New way of business thinking in TELCO business

2.1 Background

Mobile operators not only have huge needs to but also major challenges to raise profitability. Because of heavy traffic growth and the demand for new network capacity, the telco business needs new ways of managing this. For example releasing a new operating system for handheld devices can double the data demand on operators.

Communications service providers (CSPs) focus on reducing operating costs and growing revenues through the adoption of new technologies. Cellular operators, incumbent telcos, cable operators and new entrants all need to update

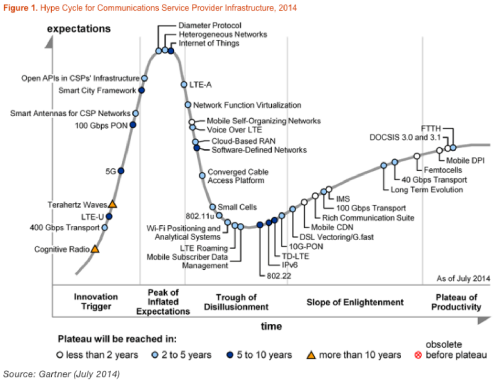

Figure 1: Hype Cycle for Communications Service Provider Infrastructure, 2014

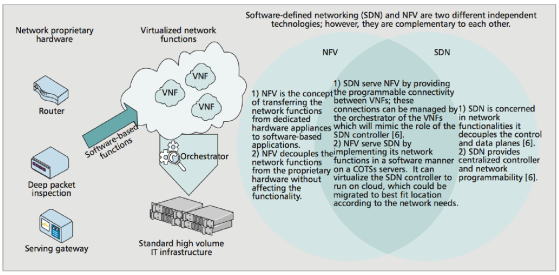

Figure 2: Network function virtualization concept

their technology adoption and retirement plans. Hype Cycle addresses the technologies needed to develop a network capable of meeting tomorrow's needs. Figure 1 shows this Hype Cycle.[17]

In order to meet good customer experience it is very important that ISP can be flexible when adjusting network and operations to satisfy their customers by responding with rapid time-to-market performance.[11]

ISP companies are going to introduce quality of service (QoS) differentiation in network. It will allow dynamic and real-time prioritization of subscribers and applications depending on network load. There will be a need to differentiate specified user groups or applications. It allows companies to build customized packages for customers and develop new revenue streams.[14]

Utilizing quality of service (QoS) functionality from operator's point of view allows operators to enhance their current business by prioritizing customers and increasing their loyalty and increasing network utilization as well as providing better visibility to the user's quality of experience (QoE).[1]

2.2 The telco cloud

Internet service providers use cloud computing model to produce telecommunication services to their customers. According to National Institute of Standards and Technology (NIST) telco cloud offers on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.

Telco cloud enables service providers to achieve network agility through software-defined networking and virtualized network functions. The telco cloud has the potential to help operators meet all their profitability and agility needs. The benefits of the telco cloud are similar to those in the IT environment, namely lower capital and operational costs and increased business agility.[11] Today, operators are in a position where they have to offer services that can exceed the boundaries of the traditional data center without compromising on quality.

New innovations are possible when adopting new capabilities into use by implementing a combination of a network-enabled cloud, which extends virtual infrastructure beyond the traditional computing and storage resourses to include network resources. Virtualization of network functionality allows portability of virtualized network functions (VNF) to different hardware platforms. This will reduce the number of platforms in the operator network.[4]

Service providers also need real-time control capabilities of software defined networking (SDN), which enables operators to more easily adapt their networks to the real-time requirements of newer services. Figure 2 shows the network function virtualization concept how NFV differs and complements SDN.[8]

Cloud-computing services are offered in the following ways: infrastructure as a service (IaaS) which is also referred to as hardware as a service (HaaS), platform as a service (PaaS), software as a service (SaaS), and network as a service (NaaS). There is no agreement on a standard definition of NaaS, but it is often considered to be provided under IaaS.

2.3 Taking NFV and new key technologies to use in the future

Virtualization of mobile networks focuses on the mobile-network base stations and mobile core network. Service providers are interested in virtualizing mobile base stations because they enable to consolidate as many network functions as possible in a standard hardware as needed to serve different mobile network technologies via a single virtualized mobile base station. LTE and 5G technology are future technologies which will use mobile network virtualization. There are many challenges in virtualizing the mobile base station. One reason is that it hosts signal processing functions in its physical layer. That is why the virtualization is first considered for implementation in the higher network stack layers. Considering eNodeB, which is the fourth-generation network (Long Term Evolution, LTE) base station, virtualization will be implemented in layer 3 and then in layer 2.[14]

Virtualizing layers 2 and 3 of the base station makes it possible to offer a centralized computing infrastructure for mul-

tiple base stations. This will lead to lower-cost base stations because only baseband signal processing should be implemented on-site. Service providers will benefit from sharing their remote base station infrastructure to achieve better area coverage with minimum CAPEX and OPEX investment.[13] There are also efforts to centralize the L1 functionalities of several base stations. They will be able to upgrade VNFs to support multiple telecommunications technologies and adapt them for new releases.[8]

In modern LTE networks we can use the Evolved Packet Core (EPC) technology as a virtualization of 4 G networks. Evolved Packet Core is a flat architecture that provides a converged voice and data networking framework to connect users on a Long-Term Evolution (LTE) network.Virtualizing the functionalities within the core is the main target for NFV. It will become essential to have a flexible and easily manageable network. A network that could be scaled on-demand in real time and be easily manageable. Virtualizing EPC offers all these benefits to service providers.[8]

In mobile networks, the radio access network as a service (RANaaS) concept has its background in the Cloud- RAN (C-RAN), NFV and mobile cloud computing concepts. The first needed component of RANaaS is C- RAN, which has centralised processing, collaborative radio, real-time cloud computing and a clean RAN system. C-RAN makes RAN more flexible and efficient. Another needed component of RANaaS is NFV, which allows network functions to run over virtualized computing environment. The European Telecommunications Standards Institute (ETSI) NFV Group allows the IT and Telecom industries to share their complementary expertise to make NFV real. It has published in details how to virtualize base station and also a framework for operating virtualized networks. RANaaS can be seen as an example of NFV, focussing on a virtualised RAN architecture that provides answers to the use of existing cloud computing concepts for the virtualisation of LTE Systems. The third needed component of RANaaS is the cloud computing concept, in other words a transition from product-based approach to a service-based approach. All cloud applications should be supplied as a service.[19]

When a mobile operator has its own LTE core network it still requires RAN access for a certain geographical area in order to serve its customers. The Mobile Cloud Network Service Provider (MCNSP) manages mobile operators subscription and is also a provider of integrated services. The mobile operator receives a service catalogue from the MCNSP and selects a RANaaS jointly provided by RAN Providers (RANPs) A and B. This allows the mobile operator to connect its customers via RANs from operators A or B to its own core network.[19]

3 Business examples

3.1 Nokia Solutions and Networks (NSN)

NSN is building its telco cloud infrastructure by using its product called Liquid Core. According to NSN's white papers, it has launched a virtualized VoLTE core which enables better multimedia experience for voice and data in the same network. An example mentioned here is cloud-enabled Voice over LTE (VoLTE). Because subscribers demand more and better services, it is necessary for operators to protect their profitability by reducing delivery costs. In NSN's view, there is solution which eliminates the need for separate voice and data networks. Voice over LTE (VoLTE) helps operators to build a next generation service architecture, which will help to ensure high quality voice, video and multimedia services for subscribers, while protecting operators' profitability.[1]

Liquid Core is a part of Nokia Telco Cloud, and it makes the core network flexible enough to meet subscribers' changing needs in order to allow operators to deliver the best customer experience at the lowest cost. Liquid Core is built on Core Virtualization and Telco Cloud Management. Using Core Virtualization enables core network functions to run on standard IT hardware in a virtualized manner. Telco Cloud Management also includes cloud-ready NetAct for managing network domains, as well as the Cloud Application Manager for managing applications in the telco cloud. Cloud Application Manager automates many phases in a telco application's lifecycle from development to deployment and later to maintenance. It also provides automated elasticity management of cloud resources to ensure that the necessary network capacity is available to run applications. Nokia Networks has demonstrated how the complete lifecycle of virtualized network elements can be managed remotely, including the deployment of virtualized VoLTE, Packet Core and Home Subscriber Server applications, while monitoring their virtual network element resources as well as their elasticity management.[7]

3.1.1 VoLTE

The Nokia VoLTE solution offers a clear evolutionary path that allows operators to provide VoLTE according to industry standards, without a need for drastic changes in the voice network. There are many business benefits: high-quality voice and video service over LTE, fast time-to-market, full mobile voice service continuity as well as core elements run on the same platform or in telco cloud by reducing total cost of ownership. It also includes design, planning and integration services.[15]

NSN is the first vendor to supply a commercial telco cloud solution which is compliant with ETSI NFV architecture for end-to-end VoLTE services. The company is also releasing Cloud Network Director, an orchestration tool that will automatically deploy, configure, optimize and repair a set of virtualized network functions to simplify the deployment of services such as VoLTE. Nokia Cloud Network Director will meet the needs of the recently published ETSI NFV Management and Orchestration specification and is the key component of fully automated network lifecycle management. Complying with the ETSI NFV Orchestrator, Cloud Network Director will have open interfaces for easy integration on Virtualized Network Function (VNF). It is integrated on existing operations and business support systems. This multi-vendor capability ensures that telco clouds can be implemented flexibly to blend in mobile broadband operators' needs and integrate with existing systems.[9]

NSN has launched a virtualized VoLTE core to enable a better multimedia experience. Virtualized VoLTE core solu-

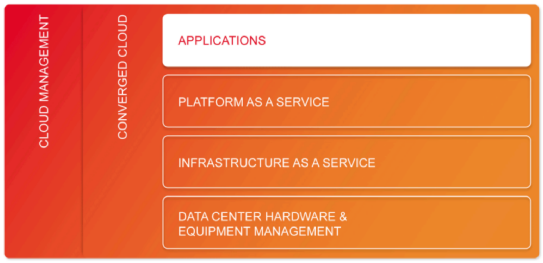

Figure 3: Ericsson Cloud System

tion consists of IMS, OpenTAS, and HSS.[1] Open TAS is a Telecommunication Application Server (TAS) in Voice over LTE (VoLTE) networks built according to IMS core architecture. Open TAS is part of Nokia Networks virtualization that provides strong Liquid Core synergies with other Nokia Networks core elements. With Open TAS operators can re-use existing back-end systems such as billing.[13]

NSN is informing a telco cloud partner certification program in order to strengthen its wide-ranging telco cloud partner ecosystem. The program enables third-party software to be certified with Nokia Networks' telco cloud solutions to deliver extra value by working with the company's virtual network functions. This approach ensures operators receive comprehensive solutions, in line with ETSI NFV, to meet the highest quality and security standards.[9]

3.2 Ericsson

Ericsson has launched its product called the Ericsson Cloud System. It is an open, distributed, telecom-grade platform based on OpenStack architecture.This cloud technology will enable NFV, making it possible to place network functions where they are best suited. Ericsson Cloud System is a open environment where new applications can run side by side with network functions.[3]

Ericsson Cloud System is a full-stack solution to handle all workloads across multiple industries. Each layer in the stack can be offered independently. They can also be combined into converged offerings, such as secure storage, providing additional value. NFV has moved into its commercial phase and it is no longer just an idea. According to NTT Docomo's press release: "a personalized mobile solutions provider for smarter living, announced today that in collaboration with three world-leading ICT vendors—Ericsson, Fujitsu and NEC—it has started developing plans for a commercial deployment of network functions virtualization (NFV) technology on DOCOMO's mobile network in Japan from March 2016, and ultimately for its entire virtualized network." In figure 3 (Ericsson Cloud System) the main offerings are presented.[2]

3.2.1 Virtual Packet Core

Ericsson is introducing its virtual Evolved Packet Core as a part of its commitment to NFV. This development of the Evolved Packet Core opens the way for new operator opportunities in many areas like M2M, enterprise and distributed cloud for rural mobile broadband. These services can be used by Ericsson Cloud System. Virtual Evolved Packet Core is made to support operators' transition to a network-enabled cloud. Virtualization is done of all Evolved Packet Core components with full compatibility. All features across native and virtualized network nodes can be used by using these functions.[5]

Virtual Evolved Packet Core includes virtual Evolved Packet Gateway, virtual SGSN-MME, virtual Service-Aware Policy Controller and virtual Service Aware Support Node. Evolved Packet Core can be run on Ericsson Cloud System using third-party hardware. It also allows virtual network functions to be aggregated to full service solutions for fast deployment.[5]

3.3 Comparison of product strategies

Both companies, Ericsson and NSN have products related to cloud computing and NFV with it. Ericsson has launched its virtual Evolved Packet Core and Nokia Liquid Core as a part of Nokia Telco Cloud respectively. Nokia's product is built on Core Virtualization and Telco Cloud Management. Core network functions can run on standard IT hardware in a virtualized manner and is referred to ETSI NFV standardization as Network Functions Virtualization.[6] Both companies have same kind of product strategies. Nokia's product supports a number of cloud stacks to meet different operators' requirements[12]. Ericsson Cloud System product is based on the OpenStack architecture [3].

Ericsson can offer data center and cloud operation services. They operate and maintain datacenters and cloud environments on behalf of their customers. Ericsson offers both customized and standardized services. Services can be located on Ericsson's or customer's datacenter.

Internet service providers can deploy a complete telco cloud infrasructure build by NSN. They can also integrate NSN telco cloud solutions of their operator clouds, based n the HP cloud.

Both companies have same vision about the future. Markets are going to change because of revolution and evolution of this technology. Because it is not possible to virtualize everything today, both companies are going to take steps towards network virtualization.

4 Discussion

Cloud computing is a new business model for ISP's and it gives a new way to offer service delivery. Operators can have better business opportunities using cloud computing and NFV. They can combine telco and internet applications from cloud for services offered to consumers and enterprises. In this cloud computing model, telco companies must think about privacy and security issues. These issues require care-

ful business and technology planning to meet all the requirements needed.[1]

The architecture's success is its modularity, well-defined functional elements and exact separation between operational and control functions. This enables operators to upgrade their network as well as quickly deploy and adapt resources to demand. All this makes it also possible for new players to easily enter the market and to permit a virtual network operator to provide connectivity to its users.[19] It is possible to deploy new services faster compared to existing methods. It is possible to achieve better hardware performance by using software based on processing in NFV. It may lead to service level agreement (SLA) violations and customer dissatisfaction.

5 Conclusion

Telco business needs new ways to reduce costs. The cloud model offers one way to do it. However, effective solutions may not be easy to find and companies must carefully think about for example security issues. NFV requires operators to find new ways of looking at basic network attributes such as performance, reliability and security. Virtualization will change many ways of configuring and managing network resources. Open source software enables developers to offer their applications and services regardless of the underlying infrastructure. Service delivery times will become shorter. Programmability, orchestration and automation are necessary components to make this happen. Various types of dedicated network appliance boxes become virtual appliances which are software entities running on a high performance server. High layer network functions become software based virtual appliances. This migration will not be easy because of many management challenges. Using NFV in data centers will be energy efficient because of fewer hardware boxes. In the future, energy efficiency will be more and more important.

Market in telecommunication business is going to change a lot in the future. NSN and Ericsson and other telecommunication companies must change their product portfolio. It will take time to migrate this step towards NFV. A lot of work needs to be done to reach full benefits of network virtualization. This migration is making constant progress.

References

[1] Change the game in content delivery, (accessed 2015-03-18). http://networks.nokia.com/portfolio/liquid-net/liquid-broadband.

[2] Ericsson cloud system, (accessed 2015-03-19). http://www.ericsson.com/spotlight/cloud.

[3] Ericsson cloud system enables network functions virtualization, (accessed 2015-02-05). http://www.ericsson.com/news/131122-ericsson-cloud-system/functions/virtualization_244129226_c.

[4] Ericsson white paper (uen 284 23-3219 rev b | february 2014). the real-time cloud, (accessed 2015-03-18). http://www.ericsson.com/res/docs/whitepapers/wp-sdn-and-cloud.pdf.

[5] Launch: Evolved packet core provided in a virtualized mode industrializes nfv, (accessed 2015-03-19). http://http://www.ericsson.com/news/1761217.

[6] Liquid core, (accessed 2015-02-05). http://networks.nokia.com/portfolio/liquidnet/liquidcore.

[7] Liquid net, (accessed 2015-03-02). http://networks.nokia.com/portfolio/liquidnet.

[8] Nfv: State of the art, challenges, and implementation in next generation mobile networks (vepc), (accessed 2015-03-18). http://ieeexplore.ieee.org.libproxy.aalto.fi/stamp/stamp.jsp?tp=&arnumber=6963800&tag=1.

[9] Nokia networks shipping industry's first commercial nfv solution networksperform, (accessed 2015-03-18). http://company.nokia.com/fi/news/press-releases/2014/09/04/nokia-networks-shipping-industrys/first-commercial-nfv/solution-networksperform.

[10] Nokia siemens networks: Cloud computing - business boost for communications industry, (accessed 2015-02-01). http://networks.nokia.com/nsn_telco_cloud_white_paper_1_.pdf.

[11] Nokia siemens networks: Cloud computing - business boost for communications industry, (accessed 2015-02-01). http://networks.nokia.com/nsn_telco_cloud_executive_summary_2014.pdf.

[12] Nsn takes telco cloud to commercial reality with new automated application management mwc14, (accessed 2015-02-05). http://networks.nokia.com/news-events/press-room/press-releases/nsn-takes-telco-cloud-to-commercial/reality-with-newautomated/application-management-mwc1.

[13] Open tas, nokia networks open telecom application server, (accessed 2015-03-18). http://networks.nokia.com/portfolio/products/ip-multimedia-subsystem-ims-core/open-tas.

[14] Quality of service differentiation, (accessed 2015-03-18). http://networks.nokia.com/portfolio/solutions/quality-of-service-differentiation.

[15] Voice over lte (volte), (accessed 2015-02-01). http://www.ericsson.com/news/150130-er-wifi-calling_244069647_c.

[16] Etsi, network function virtualization: An introduction, benefits, enablers, challenges, and call for action, (accessed 2015-03-18). http://portal.etsi.org/NFV/NFV_White_Paper.pdf, 2012.

[17] Hype cycle for communications service provider infrastructure, (accessed 2015-04-01). http://www.gartner.com/document/2806117, 2014.

[18] M. Liyanage. *Software Defined Mobile Networks (SDMN): Beyond LTE Network Architecture*, book Chapter 9. Wiley, first edition, 2014. Sakari Luukkainen, Antti Tolonen.

[19] A. H. e. a. Lucio Studer Ferreira, Dominique Pichon. An architecture to offer cloud-based radio access network as a service, (accessed 2015-03-19). http://ieeexplore.ieee.org.libproxy.aalto.fi/stamp/stamp.jsp?tp=&arnumber=6882627.

Something You Need To Know About Bluetooth Smart

Rui Yang
Student number: 467656
rui.yang@aalto.fi

Abstract

Bluetooth 4.x is the latest version in Bluetooth family. It embraces the features of BLE (Bluetooth Low Energy), which is one of the reasons contributing to Bluetooth being considered an ideal platform for the Internet of Things (IoT). Compared to previous versions of Bluetooth, BLE improves power saving, lower deployment cost and enhanced radio range, etc[5]. This article focuses on the evolutions of BLE from Bluetooth 4.0 to Bluetooth 4.2, the reasons behind these changes, an analysis of a special use case and security concerns in BLE.

KEYWORDS: Bluetooth Low Energy, BLE, Bluetooth Smart, Smart Hotel, Security

1 Introduction

Bluetooth is a wireless communication technology for short range communications especially dedicated in personal area networks (PAN). Launched about 21 years ago, it has been widely implemented and provides great conveniences to our daily life. For instance, more than 2.5 billion Bluetooth products were shipped in 2013[3]. Benefiting from the low energy consumption of BLE, a button sized cell can provide a Bluetooth instance running with singular module for more than 3 years, depending on different use cases[5], which points the way for a huge number of possible applications nowadays and in the future.

Since Bluetooth is one of technologies that form PANs, which means it is more connected to our body area, in addition to the significant importance and privacy of transmitted data, its protection of users' privacy and security issues naturally become the main focus of its users. Thus this article introduces the potential challenges regarding the concerns of privacy and security issues, and possible solutions.

2 Background

Bluetooth Low Energy is currently hosted by Bluetooth Special Interest Group (SIG), the design goals of which are lowest cost and easy of deployment. Since the classic Bluetooth is connection oriented, designed for data streaming in previous versions, it means the Bluetooth instances have to keep the connection alive even when there is no data needed to be transmitted. Lacking of sleep mode, classic Bluetooth suffers from considerable energy consumption. This is not good for some devices which only have a button cell battery. In order to overcome this shortcoming, SIG introduces Wibree of Nokia to be part of the Bluetooth standard and this standard evolved to BLE later on. It mainly focuses on keeping the energy consumption as low as possible. To achieve this, classic Bluetooth has been redesigned (not optimized from the classic Bluetooth) in BLE from radio, protocol stack, profile architecture and qualification regime[9].

Bluetooth Version	Speed
v 1.1	1Mbps
v 2.0	3Mbps
v 3.0	54Mbps
v 4.0	0.3Mbps

Table 1: Transmission Speed Over Different Bluetooth Version

We may benefit a lot from the BLE. However, in some use cases, the role of classic Bluetooth cannot be replaced by BLE. For instance, as table 1 shows, being limited by the transmission speed of BLE (mainly in Basic Rate, optional in Extend Data Rate), it would no longer be suitable for some applications which require huge amounts of data transmission, such as data streaming for music from a mobile phone to a headset. In order to be compatible with classic Bluetooth, BLE instance can run in either single-mode or dual-mode. The instance running in the single-mode cannot communicate with classic Bluetooth while the instance running in the dual-mode can.

Moreover, the BLE has the Client/Server structure in its attribute protocol. This can allow BLE devices connect to the wide area network only using a gateway, such as a PC or mobile devices. With the expansion of the concept of IoT, in addition to the ongoing integration of IPv6 into Bluetooth 4.2, it has been estimated that more than 2 billion units of BLE will be deployed around the globe[10]. The increasing popularity and huge amounts of deployment[3] requires thorough and careful analysis of potential issues associated with BLE. This is one of the reasons why this article is written to address these issues.

One of the major competitors of BLE is IEEE 802.15.4, known as ZigBee, which uses the same radio frequency as BLE. But the shipments of ZigBee instances are not comparable with BLE[1]. It mainly because ZigBee is not embedded in commonly used PCs and mobile phones, whereas BLE is. With the rapid development of IoT, the shipment of BLE instances will even be bigger. From the perspective of techniques, although ZigBee is low power and its stack is quite light, BLE has even lower power and a lighter stack[9].

This article is organized as follows. Section 3 introduces

the evolution of BLE essential features. Section 4 describes one scenario of use cases. Section 5 presents the security features of BLE. Section 6 concludes the article.

3 Evolutions of BLE Essential Features

3.1 Introduction to BLE

BLE (also known as Bluetooth Smart) was introduced in Bluetooth 4.0 specification in 2010. It enables low power devices and sensors to connect to the latest mobile devices. Compared to previous versions of Bluetooth, BLE is newly designed and has a distinct feature of low energy consumption. What's more, to make the Bluetooth devices easier to use, the discovery procedure and setup of services between devices have been simplified in Bluetooth protocol. Discovery procedure has been simplified by decreasing the number of steps in it. And simplifying setup of services is implemented by advertising all of the services which are available in one Bluetooth device to another one so that the configuration of the incoming connection can be more automated.

Normally, the design goal determines a product in respect of functionalities and performances. The goal of the classic Bluetooth is to standby for several days or to stream for several hours, while the BLE is designed to standby for several years collecting or broadcasting data such as temperature and location information. The earlier design of BLE equips with several key features, including low cost, supporting worldwide operation, low power consumption and robustness, etc. All of these design goals determine how each sub-systems in BLE should be implemented. In order to achieve the lowest cost, the system must be kept as small and efficient as possible and new methodologies should be adapt to boost the performance. For instance, to provide supports for new network topologies, BLE has been optimized to lower the cost using research based methodology[8]. In addition, BLE uses the 2.45GHz ISM band to transfer signals to support worldwide operations. However, this radio band is unlicensed and any organization can use it for commercial purpose. As a result, it is crowded with many transmission signals such as Bluetooth and Wi-Fi. In order to co-exist in such a radio band, a mechanism called Adaptive Frequency Hopping has been introduced to help Bluetooth avoid signal conflicts. Last but not least, the low energy consumption feature has had a great impact on the design of the protocol stack of BLE. For example, the link layer has been considered as the most complicated layer in the Bluetooth protocol stacks. In BLE, the link layer has been lightened and even provides ultra-low energy consumption. The main reasons behind low energy consumption are the low duty cycles and the low sleep power. As mentioned above, in classic Bluetooth, once a connection has been established, the connection should be active all the time even without data transmission. While the duty cycle in BLE is extremely low which means that the connection has a low active time and the device at other time can stay in the sleep mode. And the low sleep power consumption leads further to the super power saving in BLE.

For a technology like Bluetooth, even with critical updates, all the traditional features of classic Bluetooth should be included in the BLE as well. In order to inherit all the features from previous versions of Bluetooth at the same time, the BLE is designed to run in two modes: single-mode and dual-mode. In single-mode implementation, only the low energy protocol stack is implemented. In dual-mode implementation, the functionalities of BLE is integrated in classic Bluetooth controller[11] (see Fig.1). Furthermore, one more feature requiring mention is its change of transmission speed. Because of a top concern of low energy consumption, BLE commonly adopts the Basic Rate (BR), with a transmission rate about 0.3 Mbps while optional with Enhanced Data Rate (EDR)(see in table 1).

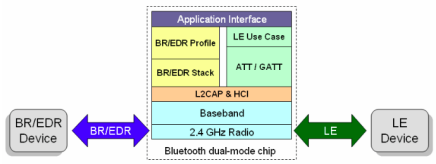

Figure 1: Dual Mode Chipsets

In short, low energy, as the critical feature of BLE, influenced BLE from its design to implementation. The number of new applications will rise in the coming future.

3.2 Evolutions introduced in Bluetooth 4.1

Bluetooth 4.1, as a critical update although without hardware component updated, has renewed its specification and been assigned more flexibility for its developers to integrate more functionalities. Moreover, a better co-existence with TLE radios, high transmission rate and high consistence of connections has also been included in Bluetooth 4.1. More specific details are described as follows.

3.2.1 Improving Usability

To improve the consumers' usabilities, which have been mentioned above, the SIG defines outcomes of these specific improvements in consumer usabilities as "just work" for a simple experience. And this "just work" experience comes from the following three aspects.

First of all, Bluetooth 4.1 provides a better co-existence with TLE (Long-Term Evolution) radios. LTE has been widely adopted as 4G standard for cellular networks and the global shipment of mobile phone supporting LTE also grows swiftly. In order to reduce the interference between these two technologies, the update in Bluetooth 4.1 allows the bluetooth device to communicate with LTE radios to minimize the interference by cooperating with each other. This is automatically done by the Bluetooth device without any operation from the user perspective.

In addition, Bluetooth 4.1 also supports the seamless and silent connection between two Bluetooth devices which have been connected with each other previously. This can help to improve the usability since the whole process is done automatically without any user's participation.

The third aspect to improve the usability is supporting bulk data transfer. The scenario is that, when large amounts

of data need to be transmitted, for instance from a sensor to a roaming device, techniques featuring data compression, data blocking and buffering have been implemented to optimize the transfer rate in Bluetooth 4.1.

3.2.2 Enabling Developer Innovation

Bluetooth 4.1 enables developer innovation in the form of allowing developers to set the role of each Bluetooth device as a Bluetooth Smart Ready Hub and Bluetooth Smart Peripherals at the same time. As a Bluetooth Smart Ready Hub, a Bluetooth device can collect data from other Bluetooth devices. While as a Bluetooth Smart Peripheral, it can transmit the data to another Bluetooth device. By setting both roles at the same time in Bluetooth devices, more use cases and applications can be built upon this innovation feature.

3.2.3 Enabling IoT

Bluetooth is a promising technology to provide wireless connectivity in the emerging world of IoT. By introducing IPv6 into Bluetooth 4.1, the device is considered as an IoT device after it connects to the public network via Bluetooth.

3.3 Evolutions introduced in Bluetooth 4.2

Published on December 3rd in 2014, Bluetooth 4.2 is a hardware update which in order to obtain the full functionalities of Bluetooth 4.2, one has to acquire a new hardware. However some of the features, including the privacy preserving, can be updated or acquired via a firmware update to the Bluetooth 4.0 and 4.1. The main updates of Bluetooth 4.2 include enhanced privacy and security, boosted transmission speed and full internet connectivity [2].

The goals of Bluetooth Security Manager are to provide security connections and secure the communications. Bluetooth 4.0 and 4.1 are vulnerable to eavesdropping and Man In the Middle (MITM) attacks. So Bluetooth 4.2 has introduced a new security model named LE security connections [7]. This new security model adopted an algorithm called Elliptic Curve Diffie-Man(ECDN) for public/private key generation. And a new key paring procedure has also been adopted in Bluetooth 4.2. It is claimed by the SIG that using the LE Secure Connections can protect the communication from passive eavesdropping and MITM attacks regardless of the paring methods [7]. More security features of Bluetooth are discussed in section 5.

Another exciting feature in Bluetooth 4.2 is its enhanced transmission rate with boosting more than 2.5 times of transmission speed and even lower transmission error [2]. And this enhanced transmission is partially achieved by increasing the payload of transmission packet. With the increase of transmission speed, the number of packets requiring to be transferred has been decreased thus fewer transmission errors could occur and less energy would be consumed.

In addition to the features mentioned above, from the aspect of enabling internet connectivities for IoT, if we consider Bluetooth 4.2 is a step into IoT which no longer needs an intermediary support to function as a gateway, such as a smartphone, Bluetooth 4.2 is another step forward with more new features released. Once the Internet Protocol Support Profile (IPSP) approves that Bluetooth can directly connect to the network using IPv6/6LoWPAN, Bluetooth can utilize the current IP structure to communicate with each other or control other devices.

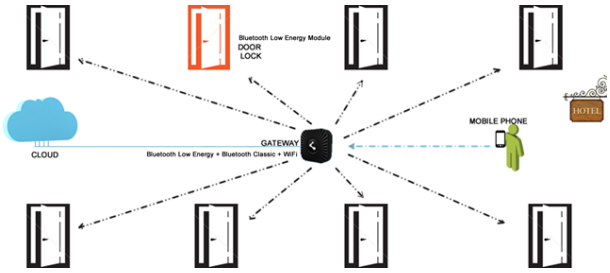

Figure 2: Components in Hotel Use Cases

4 Use Case Demonstration

With the appearance of Bluetooth, large amounts of applications based on it have been introduced to this world. This section describes one particular use case of Bluetooth in Smart Hotel. It particularly presents how the Bluetooth cooperate with other technologies and how the application is fully compatible with both old and new versions of Bluetooth. Moreover, the possible attacks and solutions have been discussed.

4.1 Use Case Introduction

Figure 2 shows all the components in the hotel use cases.
1. Mobile Devices: equipped with either classic Bluetooth or BLE;
2. Central Gateway Module: hardware which can connects to the Internet and communicate with classic Bluetooth or BLE devices;
3. Lock Modules: Physical lock embedded with BLE which controls the openness of the door and communicates with the central gateway module;
4. Cloud: backend server for authentication.

The tenant holds mobile devices and owns the credential. When they want to open their door, the mobile device connects to the central gateway module through classic Bluetooth or BLE along with user's credential. If the mobile device is using BLE, the central gateway module then also uses BLE to communicate with the mobile device. If the mobile device is not compatible with BLE, a classic Bluetooth connection will be created between them. Once the central gateway module received the credential from the mobile device, it forwards the credential to the backend server to verify the identity of the mobile device. Once the credential has been verified, the specific door will be opened for the tenant. So this is how the system works when tenant try to open the door. For such a system to work in the scenario of hotel, there are many technologies involved, cooperating with each other to fun as a whole. While what will be specifically focused is the role that Bluetooth has played in this system.

In the implementation of this system, Bluetooth plays the role of providing wireless connectivity between several mod-

ules and delivering contents between them. At beginning, the mobile phone initializes the connection with the Bluetooth device in Central Gateway module. If the credential provided by the mobile phone has passed the verification, another Bluetooth connection will be made between the Central Gateway Module and Lock Modules of one specific door. The Central Gateway Module informs the Lock Modules to open the door for the tenant. So, three Bluetooth devices have been involved in this action. The Bluetooth device acts as a Bluetooth Smart Peripheral in mobile phone within the connection with the Central Gateway Module. It transmits the data, namely the credential, and receives the response. The Bluetooth device in Lock Module acts as a Bluetooth Smart Ready Hub because it constantly listen to possible connection requests from Central Gateway Module. Moreover, the Bluetooth device in Central Gateway module acts as both the Bluetooth Smart Ready Hub and Bluetooth Smart Peripheral.

There are still numerous of applications associated with Bluetooth. This specific use case demonstrates part of abilities that the Bluetooth is capable of.

4.2 Security Concerns

The importance to keep user's credential in secret under the scenario of Smart Hotel is the same as the importance to keep the normal metal key in normal hotel use case. A disclosure of the credential would cause great damages not only to the user and reputations of the hotel but also to people's trusts on Bluetooth technology. This section mainly discusses potential risks of disclosure of credentials and possible strategies to improve security services in Smart Hotel in the usage of Bluetooth. So other potential risks caused by other part of Smart Hotel will not be discussed in this article, such as the security issues within Cloud or Central Gateway, etc.

There are two Bluetooth connections in Smart Hotel use case: the BLE connection between Central Gateway and Lock Modules, and possible BLE or classic Bluetooth connection between Central Gateway and Mobile Devices.

Firstly, in the connection between Central Gateway and Lock Modules, inappropriate settings of Bluetooth devices can lead to the disclosure of user's credential even with the latest Bluetooth version. In Bluetooth 4.0, there is no eavesdropping protection at all and the 'Just Work' paring method provides no MITM protection. Eavesdroppers can capture secret keys during paring procedure. Furthermore, MITM attack can reveal the credential transmitted between Central Gateway and Lock Modules. To minimize while not eliminate risks of eavesdropping and MITM attacks, the 'Just Work' paring method should never be used. Another risk that could disclose the credential is the improper storing of link keys. Link keys can be read and modified by attackers if they are not securely stored and protected through access controls. In addition, a weak strength of pseudo-random number generator can also lead to leak of credential. As it might appear familiar to you, during the paring procedure, user is asked input six digital which is randomly generated by another paring Bluetooth device. Since this sequence of six numbers is generated by Random Number Generator (RNG) using pseudo-random algorithm, if RNG produce static or periodic numbers, attackers can easily know this six numbers. Therefore, Bluetooth should use a RNG with strong strength of generating random number.

Concerning the Bluetooth connection between Central Gateway and Mobile Device, if it is a BLE connection between them, it is the same case mentioned above. While if it is a classic Bluetooth connection, there are also several risks needed to be mentioned in here. First of all, the repeatable attempts for authentication has the risk of reveal credential. Bluetooth specification do requires an exponentially increasing waiting interval between successive authentication attempts. While it does not require such a waiting interval for authentication challenge requests. As a result, attacker can acquire the secret link key by collecting large amount of challenge responses which are encrypted by link key. So, to minimize this risk, the mechanism which uses the same waiting interval for authentication challenge requests as successive authentication attempts. What's more, classic Bluetooth connection is also vulnerable to MITM attack. For instance, device authentication is simple shared-key challenge/response. So the countermeasure is implementing mutual authentication. In addition, The stream cider used in Bluetooth BR/EDR encryption is relatively weak. Since there are more risks in classic Bluetooth, it is highly recommended that the Mobile Devices should create a BLE instead of classic Bluetooth connection to Central Gateway Module if it is capable of.

In short, the setting of Bluetooth devices should be set as strict as possible to protect user's credential with the possibility of sacrificing a small mount of easy of usability. It is the developers' concern to determine the setting of Bluetooth connection based on use cases in their applications.

5 Security Features

When one Bluetooth device (Source) is communicating with another one (Target), other Bluetooth devices within their transmission range can also receive the signals transmitted between them. To create and maintain a secure Bluetooth connection is to make sure that either the message exchanged between them can not be intercepted, modified and interpreted by other anonymous devices or the message received comes from the Source that the Target trusts. With this kind of purpose, BLE has introduced several services to make it happen. For instance, Advanced Encryption Standard (AES) has been used to encrypt and decrypt the messages exchanged between BLE devices. Signature-based message signing method is used to provide and verify the identity of one message. Latter subsections describe more detail implementation of security services in BLE.

5.1 Encryption and Authentication services

For two "strange" BLE devices to establish a secure link between each other by using pairing, one (Initiator) initiates a paring procedure with another device (Responder) to exchange their identity for setting up trust and readying encryption keys for the data exchange [4]. During the paring procedure, they decide which paring method should be used and what to share between them. After the paring procedure,

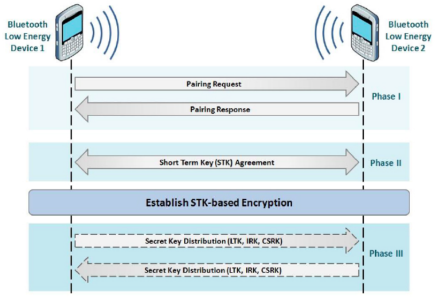

Figure 3: BLE Paring

	Initiator				
Responder	DisplayOnly	Display YesNo	Keyboard Only	NoInput NoOutput	Keyboard Display
Display Only	Just Works Unauthenticated	Just Works Unauthenticated	Passkey Entry: responder displays, initiator inputs Authenticated	Just Works Unauthenticated	Passkey Entry: responder displays, initiator inputs Authenticated
Display YesNo	Just Works Unauthenticated	Just Works (For LE Legacy Pairing) Unauthenticated	Passkey Entry: responder displays, initiator inputs Authenticated	Just Works Unauthenticated	Passkey Entry (For LE Legacy Pairing): responder displays, initiator inputs Authenticated
		Numeric Comparison (For LE Secure Connections) Authenticated			Numeric Comparison (For LE Secure Connections) Authenticated

Figure 4: Paring methods part one

all the encryption keys should have been exchanged and thus creating and maintaining a secure connection. The most critical step among all these three steps is what happened during paring procedure and what secure services mainly reveal in this round. The paring procedure is consists of three phases (see Fig.3).

The first phase is two devices announcing their input and output capabilities and choosing a suitable method for phase two [6] (see picture 4 and 5).

The second phase is to make sure that the key exchange in phase three is operated in a secure way by generating a Short-Term Key (STK) to encrypt the transmission channel. The two devices agree on a Temporary Key (TK) using the method chosen in the first phase. One common method that the reader might be familiar with is using the Passkey Entry method, in which the user is asked to input six random digits as the TK. Another one is assisted by other technologies (such as NFC) for the TK agreement, which is called the Out of Band method. If both methods are not available, the final one, called the Just Works method is adopted. This method has no authentication at all and thus lacks of the ability to prevent a MITM attack. After getting the TK, the STK can be obtained for each paring devices using this algorithm: STK = AES128 (TK, Srand || Mrand).

After the STK is obtained, up to three 128-bit keys can be encrypted using STK and distributed to another device. These three keys contains Long-Term Key (LTK), the Connection Signature Resolving Key (CSRK) and the Identity Resolving Key (IRK). LTK is used to create a session key

	Initiator				
Responder	DisplayOnly	Display YesNo	Keyboard Only	NoInput NoOutput	Keyboard Display
Keyboard Only	Passkey Entry: initiator displays, responder inputs Authenticated	Passkey Entry: initiator displays, responder inputs Authenticated	Passkey Entry: initiator and responder inputs Authenticated	Just Works Unauthenticated	Passkey Entry: initiator displays, responder inputs Authenticated
NoInput NoOutput	Just Works Unauthenticated	Just Works Unauthenticated	Just Works Unauthenticated	Just Works Unauthenticated	Just Works Unauthenticated
Keyboard Display	Passkey Entry: initiator displays, responder inputs Authenticated	Passkey Entry (For LE Legacy Pairing): initiator displays, responder inputs Authenticated	Passkey Entry: responder displays, initiator inputs Authenticated	Just Works Unauthenticated	Passkey Entry (For LE Legacy Pairing): initiator displays, responder inputs Authenticated
		Numeric Comparison (For LE Secure Connections) Authenticated			Numeric Comparison (For LE Secure Connections) Authenticated

Figure 5: Paring methods part two

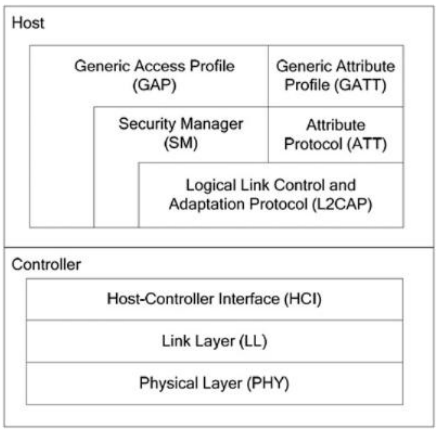

Figure 6: Bluetooth Blow Energy Protocol Stack

for each Link Layer connection (See BLE protocol structure in Fig. 6). CSRK is to sign the data at the ATT Layer and IRK is to generate the private address using the known public address thus avoiding the tracking.

In short, this is the whole three phases that enable both BLE 4.0 and 4.1 devices to build trusts between each other. One possible threat about this mechanism is introduced in the following section as well as the corresponding solution adopted in BLE 4.2.

5.2 Possible Security Issues and Solutions

This section mainly describes part of potential threats to the BLE devices. As mentioned in previous section, the BLE 4.0 and 4.1 is vulnerable to MITM attack. The most common is the active eavesdropping in which the hacker acts as the intermediary between two BLE victim devices, while the victims believe that they are directly connected to each other. Thus the hacker can acquire all the data that have been transmitted between these two. This attack is also common in the HTTPS web applications. To prevent the active eavesdropping, two methods, namely Passkey Entry and Out of Band method, can be used.

Another critical MITM attack is passive eavesdropping, in which the hacker silently listens to the connection without victims being aware of his presence. This can happen when the keys (in 3 phase mentioned above) are exchanged through an insecure channel and the hacker intercepts these keys to decrypt the messages it receives from the victims. In order to overcome this issue, Elliptic curve Diffie Hellman (ECDH) has been introduced into LE Secure Connection.

6 Conclusion

This article mainly introduces essential features of Bluetooth 4.0 to 4.2 along with their security features. It introduces the essential features from 4.0 to 4.2 at the beginning. Then one specific use case in Smart Hotel was discussed and the roles that Bluetooth played in such a use case have been described as well as potential risks of disclosure of credentials and advices to minimize these risks. After that, security features of BLE were introduced, along with possible threats and solutions presented.

References

[1] Z. Alliance. Market Leadership. Technical report, ZigBee Alliance, Available on February 2014. http://old.zigbee.org/About/AboutTechnology/MarketLeadership.aspx.

[2] Bluetooth.org. New Bluetooth Specifications Enable IP Connectivity and Deliver Industry-leading Privacy and Increased Speed. Technical report, Bluetooth Special Interest Group, 2014. http://www.bluetooth.com/Pages/Press-Releases-Detail.aspx?ItemID=220.

[3] Bluetooth.org. History of the Bluetooth Special Interest Group. Technical report, Bluetooth Special Interest Group, Available on February 2015. http://www.bluetooth.com/Pages/History-of-Bluetooth.aspx.

[4] Bluetooth.org. Security, Bluetooth Smart (Low Energy). Technical report, Bluetooth Special Interest Group, 2015. https://developer.bluetooth.org/TechnologyOverview/Pages/LE-Security.aspx.

[5] Bluetooth.org. The Low Energy Technology Behind Bluetooth Smart. Technical report, Bluetooth Special Interest Group, Available on February 2015. http://www.bluetooth.com/Pages/low-energy-tech-info.aspx.

[6] J. P. Carles Gomez, Joaquim Oller. Overview and Evaluation of Bluetooth Low Energy: An Emerging Low-Power Wireless Technology. Technical report, Universitat PolitÃĺcnica de Catalunya/FundaciÃş i2Cat, August 2012. http://www.mdpi.com/1424-8220/12/9/11734/htm.

[7] V. Gao. Everything that you always want to know about Bluetooth security in Bluetooth 4.1. Technical report, Bluetooth Special Interest Group, 2014. http://blog.bluetooth.com/everything-you-always-wanted-to-know-about-bluetooth-security-in-bluetooth-4-2/.

[8] R. Heydon. *Bluetooth Low Energy: the developer's handbook*. 2013.

[9] S. A. Joe Decuir. Bluetooth 4.0: Low Energy. Technical report, CSR plc, Available on February 2015. http://chapters.comsoc.org/vancouver/BTLER3.pdf.

[10] A. West. Smartphone, the key for Bluetooth low energy technology. Technical report, IMS Research, Available on February 2015. http://www.bluetooth.com/Pages/Smartphones.aspx.

[11] Wikipedia. Bluetooth. Technical report, avaiable on Feb. 16th, 2015. http://en.wikipedia.org/wiki/Bluetooth.

A Survey of Password Managers

Can Zhu
Student number: 467711
can.zhu@aalto.fi

Abstract

Password management is a high profile topic these days because users employ different passwords every day. At the same time, a number of password managers (PSMs) have appeared to help users manage their passwords. Some of password managers are browser-based, and some of them are device-based. This survey analyzes three popular browser-based PSWs (Google Chrome, Firefox, Internet Explorer) and one device-based PSM (KeePass) and evaluates their security. There is also a comparison among these four password managers after analyzing their security problems. We find that KeePass is the most secure one, followed by Firefox, then Internet Explorer and Google Chrome.

KEYWORDS: password manager(s), brower-based, device-based, security

1 Introduction

The most common technique for user authentication is the use of passwords. However, how to manage passwords is among the most vexing issues in systems administration and computer security. As password management becoming more and more important, a large number of research contributions have been made toward increasing the security and usability of password-based authentication[3]. This paper classifies four password managers into two categories, browser-based and device-based, analyzes their working processes and evaluates their security. Also, it makes a comparison among those managers.

In this survey, Sec. 2 gives the background. Browser-based and device-based password manager are presented in details in the next two sections Sec. 3, Sec. 4. Finally, Sec. 5 concludes the paper and Sec. 6 points out the future work.

2 Background

Most websites, applications, and mobile APPs use passwords to authenticate their users, but they have many usability problems and weaknesses. Password security depends on choosing passwords that are unique and hard to guess, yet it is difficult to remember and retype these passwords correctly. "The passwords that are easiest to choose and memorize tend to be vulnerable to *dictionary attacks*, in which an attack tries to guess the password by constructing likely possibilities from lists of words and common passwords" [16]. As for users, they have to memorize correspondent passwords when logging into the websites, email accounts, Android applications, etc. Moreover, users need specific passwords for each account. If they just use the same passwords everywhere and someone finds this password, they will have a serious problem. The thief would be able to access user's email account, bank account, etc. However, it is nearly impossible to remember a such number of passwords, thus many users allow password managers to do the task for them. Changing passwords frequently helps to resist attack, but this makes memeory of passwords even harder.

Object	Master password
Google Chrome(41.0)	No
Internet Explorer(9.0)	No
Firefox(15.0)	Yes
KeePass(2.28)	Yes

Table 1: A table with survey objects features

Under this circumstance, a variety of password managers (PSMs) appear. Types of services they provide include password strengthening through iterated hashing[6, 4, 13], phishing protection[13], and converting other types of authentication into password[1]. Another main service provided by PSMs is the stroage and retrieval function. Some of PSMs are browser-based and some of them are device-based. This paper focuses on three browser-based PSMs, namely Google Chrome, Internet Explorer and Firefox and one device-based PSM called KeePass. And describes their working process and analyses their security. Table 1 shows PSMs studied in this paper and gives their master password features. In a password manager, when a user name is given, the system can fill the matched password automatically. A simple management system stores password, whereas a more sophisticated system locks the password under a master password. For example, the KeePass helps a user to manage his password by putting it in a database. It is locked by the master password, and user only needs to remember one master password to unlock the whole database. Also, the complete database is encrypted by the most secure encryption algorithms currently known, not only the password field.

3 Browser-based password managers

The popular Web browsers Google Chrome, Internet Explorer and Firefox all provide users with password managers. Firefox's PSM has a master password, while Google Chrome and Internet Explorer PSM does not. R. Zhao and C. Yue at [17] give an interesting analogy between the password man-

ager of a browser and its master password. They describe a scenario where values in home are stored in a safe. Because nobody guarantee that a thief cannot enter the home, a solution to open the safe cannot be early found. A computer is similar to a home, a PSM is similar to a safe and a master password is similar to the code to the safe. Though nobody could prevent attackers breaking into the computer, the possibility of decrypting the passwords via using master password should be decreased. This section discusses how three PSMs manage passwords and what security problems they have faced.

3.1 Working Process

The *automatic autofill* function of Google Chrome, Internet Explorer and Firefox are able to recongnise usernames and password fields as soon as the login page is loaded without requiring any user interaction[14].

Google Chrome and Internet Explorer perform encryption and decryption under Windows 7 with the help of two Windows API functions respectively, namely *CryptProtectData* and *CryptUnprotectData*. Fig. 1 shows how to manage passwords in Google Chrome and Internet Explorer. The benefit of using these two functions is that the encrypted text can only be decrypted under the same Windows user account. Google Chrome does not provide additional entropy to these two API functions. It just stores each plaintext username, encrypted password, and plaintext login webpage URL into the logins table of an SQLite database file named *Login Data*[17, 14]. Because PSM allows the password to be autofilled on a page within the same domain as the page from which the password was originally saved, the next time users visit the URL addresses, their login information is autofilled by Google Chrome. "Internet Explorer encrypts each username and password pair and saves the ciphertext as a value data under the Windows registry entry. Each saved value data can be indexed by a value name, which is caculated by hashing the login webpage URL address"[17]. Different from Google Chrome, Internet Explorer provides the login webpage URL address as the additional entropy to the two API functions.

Though Firefox allows users to set a master password to further protect their information, it is not enabled by default. If users set their own master password, it will be more secure for their personal information. The Google Chrome directly takes user names, passwords as plaintext before encrypting. Compared to that, the Firefox uses a master password to encrypt the original user name and password before further encrypting them. Thus, once the user name and password are obtained by attackers, they have to decrypt the encrypted text. Fig. 2 shows how to manage passwords in Firefox.

3.2 Security challenges

The biggest problem of Google Chrome is that the password can be got easily when an attacker enters users' computer. For example, if the attacker goes to the browser's settings and clicks on the show button in the preferences tab, he can obtain any saved password (Fig. 3). Internet Explorer is more secure since it does not allow some one view

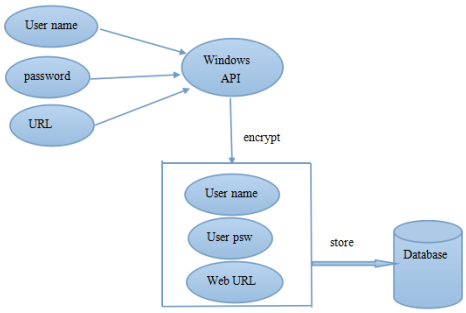

Figure 1: Password Manegement in Google Chrome and Internet Explorer

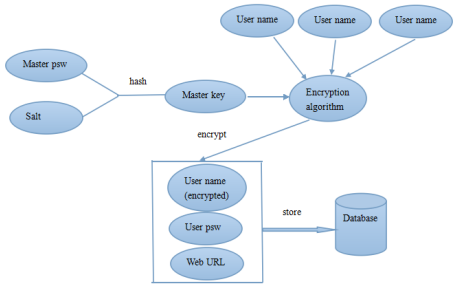

Figure 2: Password Manegement in Firefox

saved password, and it aslo does not synchronize users' data across computer. Chrome and Explorer take the user's computer login passwords as the cipher for the encrypted data, thus, it is easy for these passwords to be revealed with tools such as WebBrowserPassView of Microsoft[12]. However, WebBrowserPassview cannot retrieve passwords that are encrypted with a master password, which makes Firefox the most secure one among these three browsers.

Even though decrypting a user's login information becomes harder, brute force attacks and phishing attacks on the master master password are still quite possible [17]. Additionally, the master password is not enabled by default. If Firefox master password is not set, it is as vulnerable as Google Chrome and Internet Explorer PSMs.

Another threat model is that hackers can temporarily install malware such as Trojan horses and bots on a user's computer using attackers such as drive-by downloads [8, 15]. By using drive-by downloads, an attacker can deliver specific decryption tools to users' computers and trigger their execution. Suddenly, all the information (user names and passwords) can be completed decrypted and sent to attacker's computer.

Google Chrome, Internet Explorer and Firefox have an *automatic autofill* function, which exposes them to other vulnerable situations. Many websites serve a login page over HTTP, but submit the users' password over HTTPS. For example, let us imagine that Bob uses a password manager to save his passwords for this sites. At some point later, Bob connects to a rogue WiFi router at a shopping center. His browser is directed to a landing page and asks him to agree

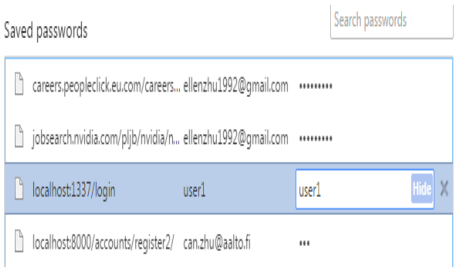

Figure 3: Gaining passwords in Chrome's settings

the terms of the service, as the condition to connect the WiFi hotpots. However, Bob is unaware that the landing page contains multiple invisible iFrames pointing to the login pages of the websites where Bob saves his sensitive information. When the browser loads these iFrams, the rogue router injects JavaScript into each page and extracts Bob's passwords autofilled by the password manager.

In JavaScript, an action attribute of a form specifies where the form's contents will be sent to after submission.
<form action="example.com" method="post" >
One way for an attacker to steal a user's passwords is to if redirect the user to another website and gain the passwords via autofill fuction. Thus, this function of the browser-based password manager increases its vulnerability.

Figure 4: The master password UI in KeePass

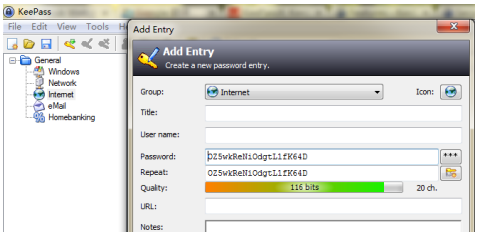

Figure 5: KeePass generates a complicated password for user

4 Device-based password managers

KeePass can be carried on a USB stick and runs on Windows system without being installed; it also can be used on Android, Linux, Mac OX and so on. Thus we include it into device-based password managers. KeePass packages are available, too. Developer of KeePass claim that "KeePass doesn't store anything on your system. The program doesn't create any new registry keys and it doesn't create any initiablization files in your Windows directory. Deleting the KeePass directory to using the uninstaller leaves no trace of KeePass on your system" [10]. Due to those features, it has become a popular password manager. This section introduces how KeePass protects users' sensitive information in detail.

4.1 Working Process

How to use KeePass on Windows As we have showed in Table 1, what improves KeePass' security is that it uses a master password to further encrypt sensitive information. Before users create a file to keep all passwords, they have to enter a master password Fig. 4. To manage the password, there are two ways. One way is that users choose the passwords themselves and ask KeePass keep them. Another way is to let KeePass generate a password Fig. 5, which is rather complicated, and save it automatically for the user. When the user needs to fill out the user and password field, he just needs to click the button to obtain them through the database on KeePass.

Encryption technology in KeePass KeePass includes all the sensitive information, e.g. master passwords, user names, user password into a database, which is encrypted completely, not only the user's password. To encrypt the database, either Advanced Encryption Standard (AES) or Twofish block cipher is applied. These two algorithms are well-known and both with high level of security. For example, AES is a U.S. Federal government standard and applies in the National Security Agency (NSA) for top secret information[5]. For both algorithms, a 128-bit initialization vector (IV) is generated randomly each time the user saves the database. "This allows multiple databases to be encryted using the same key without observable patterns being revealed"[11].

In KeePass, the user can choose both password and key file to be authenticated. The Secure Hash Algorithm *SHA-256* is used to generate the 256-bit key for the block ciphers. This algorithm compresses the user key provided by the user (consisting of password and/or key file) to a fixed-size key of 256 bits. If only a password is used, it and a 128-bit random salt are hashed using *SHA-256* to form the final key. If both password and key file are used, the final key is derived as follows: *SHA-256* (SHA-256 (password), key file contents) or *HSA-256* (SHA-256 (password), HSA-256 (key file contents))[11]. It is impossible to invert the process currently, which means one cannot compute the hash function or find a second message compressed as the same hash.

Besides using strong encryption e.g., AES-256 to protect the password database, KeePass saves the hash of the master key and entry passwords in process memory at run-

time. ”Specifically, within 21 microseconds after the entry passwords have been read and decryted from the password database”[7]. By using this technology, attackers will not be able to find any sensetive data even though they can dump the KeePass process memory to their disk.

4.2 Security chanlleges

KeePass requires some user interaction before autofilling, for instance, clicking or typing into the user name field, pressing a keybroad shortcut, or pressing a button in the browser, etc. This requirement makes it more secure than automatic autofill PSMs in some degree. However, KeePass still faces some security problems, that are described below.

Using the master password mechnism can better protect the passwords, but it should be carefully designed to maximize the security. One main security concern is a brute force attack on the master password. If the computation time for verifying a master password is very short in KeePass, it is still possible to effectively perform brute force attacks on a user's master password[17].

To increase the security, KeePass saves the hash of the master key and entry passwords in process memory at runtime. However, a framework named CipherXRay can recover the entry passwords and the master key from KeePass by determing exactly when and where the entry passwords and the master key will be briefly unencrypted in the memory[7]. In a test, X. Li, X. Wang, and W. Chang used KeePassX to create a 1468-byte encrypted password database of four entries. After tainting the encrypted password database, they used CipherXRay to monitor the execution of KeePassX. CipherXRay successfully identified that there was a 128-bit block cipher decryption in CBC mode of operation on the data read from the database file, and further identified the 256-bit secret key that was derived from the master password. Finally, CipherXRay also successfully recovered the four entry passwords in palintext form. After that, these scientists used the OpenSSL command line tool to collected information and obtained all the four plaintext password entries.

5 Conclusion

In Windows 7, Google Chrome and Internet Explorer use the Windows API functions *CryptProtectData* and *CryptUnprotectData* to encrypt and decrypt sesentive information respectively. Without a master password, they saves each plaintext username, encrypted password, and plaintext login webpage URL address into the logins table of SQLite database files. Google Chrome does not provide any additional entropy to the two API functions, while Internet Explorer does.

Firefox saves each encrypted username, encrypted password, and plaintext login webpage URL address into the login table of an SQLite database file named signons.sqlite[17]. Under a master password, Firefox hashes it with a global 160-bit salt using a SHA-1 algorithm to generate a master key. This key will be used to encrypt three Triple-DES keys, a string ”password-check” and saves the ciphertext to the key3.db file; Later, Firefox will decrypt this ciphertext to authenticate a user.

Besides using a master password mechanism, KeePass also has a protection against dictionary attacks. To generate the final 256-bit key used for the block cipher, KeePass first hashes the user's password using *SHA-256*, encrypts the result N times using the AES algorithm, and then hashes it again also by using a 256-bit key. For AES, a random 256-bit key is used, which is stored in the database file. As the AES transformations are not precomputable, an attacker has to perform all the encryptions, too, otherwise he cannot try and see if the current key is correct. The attacker now needs much more time to try a key. If he can only try a few keys per second, a dictionary attack is no longer practical. Though it is impossible to prevent dictionary attacks totally, KeePass makes this type of attack harder via this feature.

Thus, the security level among these four PSMs is:
KeePass>Firefox>Internet Explorer>Google Chrome

Some services provided by managers to strengthen password include master password, dictionary attacks protection, etc. Major browsers e.g., Firefox, Chrome, offer a built-in password wallet. However, they still suffer from security problems. Maybe these are the reasons why experts refuse to use such password managers, prefering instead to manage their passwords via writing them down somewhere nearly safe relying on memory.

6 Future work

As it mentioned above, a sophisticated password manager with a master password has a higher security level than a naive one. But this protection allows offline attacks on the master password[2]. Thus, we try to find a password manager which is even more secure than that. For example, McCarney, Daniel and Barrera, David and Clark, Jeremy and Chiasson, Sonia and van Oorschot, Paul C. in [9] present a type of password manager called **Tapas**, which combine usability and security, and also easy to implement. In this manager, passwords are protected aganist offline attacks with a strong encryption key. Using of this type of manager requires control of two independent devices, without master password. **Tapas** is a smartphone-assisted password manager for a computer. It encrypts and stores the passwords on a smartphone, and keep the decryption key inside the browser on the paired computer. If one of the two devices is stolen, the thief cannot recover the passwords in practice. This is called *Dual-possession authentication*, in which involves two applications, a *manager* and a *wallet*. They are on different devices and offer the three protocols to manage the passwords: Pair(Protocol1), Store(Protocol2), and Retrieve(Protocol3). "These protocols are designed to achieve a relatively simple goal: by stealing the data of either the Manager or the Wallet, an adversary cannot determine the stored password for any given account with any greater success than attacking the account directly" [9].

In the future, we intend on continuing the explore of password managers, and try to find more effective PSMs.

References

[1] R. Biddle, S. Chiasson, and P. C. Van Oorschot. Graphical passwords: Learning from the first twelve years. *ACM Computing Surveys (CSUR)*, 44(4):19, 2012.

[2] J. Bonneau. The science of guessing: analyzing an anonymized corpus of 70 million passwords. In *Security and Privacy (SP), 2012 IEEE Symposium on*, pages 538–552. IEEE, 2012.

[3] J. Bonneau, C. Herley, P. C. Van Oorschot, and F. Stajano. The quest to replace passwords: A framework for comparative evaluation of web authentication schemes. In *Security and Privacy (SP), 2012 IEEE Symposium on*, pages 553–567. IEEE, 2012.

[4] X. Boyen. Halting password puzzles. In *Proc. Usenix Security*, 2007.

[5] W. B. M. D. J. F. E. R. Elaine Barker, Lawrence Bassham. Aes comments by nist, March 7th, 2015. `http://csrc.nist.gov/archive/aes/round2/r2report.pdf`.

[6] J. A. Halderman, B. Waters, and E. W. Felten. A convenient method for securely managing passwords. In *Proceedings of the 14th international conference on World Wide Web*, pages 471–479. ACM, 2005.

[7] X. Li, X. Wang, and W. Chang. Cipherxray: Exposing cryptographic operations and transient secrets from monitored binary execution. *Dependable and Secure Computing, IEEE Transactions on*, 11(2):101–114, 2014.

[8] N. P. P. Mavrommatis and M. A. R. F. Monrose. All your iframes point to us. In *USENIX Security Symposium*, pages 1–16, 2008.

[9] D. McCarney, D. Barrera, J. Clark, S. Chiasson, and P. C. van Oorschot. Tapas: Design, implementation, and usability evaluation of a password manager. In *Proceedings of the 28th Annual Computer Security Applications Conference*, ACSAC '12, pages 89–98, New York, NY, USA, 2012. ACM.

[10] K. Organization. Keepass features, March 2nd, 2015. `http://keepass.info/features.html`.

[11] K. organization. Keepass help center, March 3rd, 2015. `http://keepass.info/help/base/security.html#secref`.

[12] M. Pinola. lifehacker: Which password manager is the most secure?, March 9th, 2015. `http://lifehacker.com/5944969/which-password-manager-is-the-most-secure`.

[13] B. Ross, C. Jackson, N. Miyake, D. Boneh, and J. C. Mitchell. Stronger password authentication using browser extensions. In *Usenix security*, pages 17–32. Baltimore, MD, USA, 2005.

[14] D. Silver, S. Jana, E. Chen, C. Jackson, and D. Boneh. Password managers: Attacks and defenses. In *Proceedings of the 23rd Usenix Security Symposium*, 2014.

[15] Y.-M. Wang, D. Beck, X. Jiang, R. Roussev, C. Verbowski, S. Chen, and S. King. Automated web patrol with strider honeymonkeys. In *Proceedings of the 2006 Network and Distributed System Security Symposium*, pages 35–49, 2006.

[16] K.-P. Yee and K. Sitaker. Passpet: convenient password management and phishing protection. In *Proceedings of the second symposium on Usable privacy and security*, pages 32–43. ACM, 2006.

[17] R. Zhao and C. Yue. All your browser-saved passwords could belong to us: a security analysis and a cloud-based new design. In *Proceedings of the third ACM conference on Data and application security and privacy*, pages 333–340. ACM, 2013.