Smart Cities
An Architectural Approach

André Duarte[1,2], Carlos Oliveira[2] and Jorge Bernardino[1,3]
[1]*Instituto Superior de Engenharia de Coimbra, Polytechnic of Coimbra, Coimbra, Portugal*
[2]*Ubiwhere, Lda., Coimbra, Portugal*
[3]*CISUC – Centre for Informatics and Systems of the University of Coimbra, Coimbra, Portugal*

Keywords: Smart Cities, Machine to Machine (M2M), Machine Learning, Internet of Things (IoT).

Abstract: Smart cities are usually defined as modern cities with smooth information processes, facilitation mechanisms for creativity and innovativeness, and smart and sustainable solutions promoted through service platforms. With the objective of improving citizen's quality of life and quickly and efficiently make informed decisions, authorities try to monitor all information of city systems. Smart cities provide the integration of all systems in the city via a centralized command centre, which provides a holistic view of it. As smart cities emerge, old systems already in place are trying to evolve to become smarter, although these systems have many specific needs that need to be attended. With the intent to suit the needs of specific systems the focus of this work is to gather viable information that leads to analyse and, present solutions to address their current shortcomings. In order to understand the most scalable, adaptable and interoperable architecture for the problem, existing architectures will be analysed as well as the algorithms that make them work. To this end, we propose a new architecture to smart cities.

1 INTRODUCTION

Nowadays most people live in urban areas. As populations grow, they place increasing demand on the city ecosystem and directly affect the entities responsible for the city control. These challenges make leaders adopt ways to engage with the surroundings of their city, making them more prepared and aware. The decisions they make not only directly affect the city in a short term, but are also a means to improve the decision making process. With the growth of human beings in urban areas comes a significant growth in data. This data comes from sensor networks scattered around the city or from the sensors in a smartphone. As data was produced there seemed to be a constant need to integrate all of this data to provide services, therefore, smart cities materialised.

There is a wide variety of city conceptions that have built a new horizon for cities in their challenging tasks in an increasingly cost-consciousness, competitive and environmentally oriented setting. Irrespective of whether the concept is smart city, intelligent city, sustainable city, knowledge city, creative city, innovative city, ubiquitous city, digital city or city 2.0 (e.g. Komninos, 2002; Aurigi, 2005; Carillo, 2006; Hollands, 2008, 305) they all paint a picture of a modern city with smooth information processes, facilitation mechanisms for creativity and innovativeness, and smart and sustainable service solutions and platforms (Anttiroiko et al., 2014). However, there is still a general absence of joint planning by city governments with utility providers (e.g. water, in respect of environmental sustainability) and other public services (e.g. health care). Cultural barriers include commercial confidentiality, whereas social media user groups work with open data systems, causing problems for joint working of cities with the private sector. This may create problems for collaborative ventures between city governments and businesses, and even with other public sector agencies, as well as with voluntary and community organisations. According to (Alazawi et al., 2014) a smart city depends on the provision of information, communication technologies and services to the population via web based services. However, the concept of smart city can, many times, be mistaken. In order to be smart, a city does not need state of the art technology, what it needs is interoperability between various key aspects of the city, such as governance, finance, transportation and many others.

Duarte A., Oliveira C. and Bernardino J..
Smart Cities - An Architectural Approach.
DOI: 10.5220/0005463105630573
In *Proceedings of the 17th International Conference on Enterprise Information Systems* (ICEIS-2015), pages 563-573
ISBN: 978-989-758-097-0

The kind of changes that smart cities will bring to the current world are many times said to be as similar to those seen in the industrial revolution. The motivation behind the concept is the ability to improve the city ecosystem while focusing on people, allowing technology to work for them and not with them, this will result in a greater vision of society.

Furthermore this data brings many possibilities to the cities because it makes smart systems' proliferation possible. One of these cases can the smart emergency management system, which is an extremely important piece for the welfare and wellbeing of people. According to (Feng and Lee, 2010) emergency management is a dynamic and continuous process that involves preparing for disaster before it happens. If these systems are in place the probability of anticipating man-made or natural disasters increases.

The systems already in place are decentralized, which means that they do not communicate between each other, making it almost impossible to prevent disaster. This decentralization is due to the objectives of the development. Most of the times these systems are designed to address a specific case or to work as an independent system that may receive information from many parts, although without the aim to deliver information to the necessary parties.

With the intent to address these shortcomings our work will provide an architecture to a smart system in the context of smart cities. This architecture will be created with awareness of the system's possibility to scale and to adapt itself to different contexts. This architecture will address the problem of receiving the data, process it and then retrieve useful outputs to any party that subscribes to a specific type of content. This architecture can then be tuned to fit different use cases and scenarios.

The remainder of this paper is structured as follows. Section 2 presents related work on the topic and aims to cover as much information as possible; Section 3 is discusses related technologies and intends to cover technological key aspects regarding the theme; Section 4 shows functional use cases with the objective of creating a baseline to support some of the decisions made during the work; Section 5 presents an architecture for future practical application of the analysed concepts and serves to document it. Finally, section 6 presents our main conclusions and suggests future work.

2 RELATED WORK

The problem presented in this paper, has been partially developed in the past years with other studies and projects. This section provides the necessary background to understand the basis of the developed work. It is important to acknowledge that the documented analysis in the paper will be high-level, in spite of covering as much information as possible.

There are many papers that present solutions for the issue that we are working on. The rest of this section will address part of them, which we think to be the best fit for our work.

In (Vakali et al., 2014) the concept of smart city is discussed due to its current vagueness. Still, according to (Vakali et al., 2014) this concept can vary from the technologies and infrastructures of a city to an indicator that measures the education level of its inhabitants. Furthermore the work intends to analyse the SEN2SOC experiment for its impact in the current context of this topic. The SEN2SOC (SENsor to SOCial) experiment promotes interactions between sensors and social networks to enhance the quality of data in SmartSantander.

The concept of smart city is also referred and conceptualized in (Chourabi et al., 2012). The work intends to create a framework that will sketch practical implications for governments. Furthermore the work enlists some success factors for smart cities, which are: (1) management and organization; (2) technology; (3) governance; (4) policy context; (5) people and communities; (6) economy; (7) built infrastructure; (8) natural environment. The proposed framework will provide integration to all of these factors and explain correlations between them.

Although the smart city concept began to be defined in the previous work, more recent works seem to extend this concept and provide different definitions for it.

In (Piro et al., 2014) discuss that there is yet to exist a theoretical definition of Smart City, although cities are developing and shaping for not so distant future. Furthermore the work enlists some of the current definitions for the concept, that there is yet to be completely defined.

The work also enlightens the necessity of Information and Communication Technology (ICT) services, with the intent to integrate them in a generic scenario of a smart city. The approach is from a service point of view, which means that it emphasises the role of the services in the city. It is also important to refer that real world cases are shown to prove the importance of the topic.

Alongside with smart cities there are many other concepts that need to be addressed, one of them is the Internet of Things (IoT).

According to (Jara et al., 2014) this concept comprises the full ecosystem of data in smart cities, which in other words means that IoT generates massive amounts of data that need to be processed by algorithms and tools with the intent to be useful for a city. This will also provide new ways to interact with intelligent devices and create homogeneous platforms that include both machines and humans working together.

Still according to (Jara et al., 2014) this new paradigm will shape the world and create a new conception of the Internet and how people interact with it, due to the constant interconnectivity between people and the world. It will also provide the necessary resources for the creation of new applications and data driven platforms that will, hopefully, improve the citizen's quality of life.

This new way of reinventing the Internet will not only provide endless possibilities to improve the overall interaction between humans and machines but also create new challenges, which need to be tackled, to cities themselves.

Furthermore, the work aims to develop data-driven models based on human actions to act as proof of concept for Smart Cities. The system was developed using the SmartSantander testbed, which contains real-time systems and sensors scattered around the city.

Additionally the work concludes that the devices in the Internet of Things are able to gather data and provide knowledge and that a new age of interaction is about to appear, due to the increasing demand for smart applications.

In (Benkhelifa et al., 2014) the authors listed the current disaster management projects. The purpose of this work is to summarize existing projects regarding this matter. This work is relevant due to its diversity and detail while presenting the projects, it is extremely important to have a baseline of what was already studied and how it can, if possible, be improved. It is important to state that the focus of this work is wireless sensor networks. The most relevant outputs of this work in this context were the knowledge and awareness of the projects in this area. This listing provided a wider perspective about the topic and led to discoveries regarding the State of the Art projects, which by itself ignited the discovery of solutions and use cases for each problem.

One of the major problems encountered when dealing with large amounts of data is the system's scalability. In order to understand how similar systems operate when larger amounts of data are in place (Albtoush et al., 2011) explains implementation choices that should be made in order to avoid problems. This provides useful outputs for the viability and feasibility of the system. This work also explains the necessity of risk assessment of the system, not only during the implementation but also during the working phase. Finally it is also important because it defines a framework for emergency management, which includes risk assessment and disaster prevention in a multilevel and multidimensional architecture.

With the intent of presenting the role of today's technologies in this field in (Alazawi et al., 2014) it is stated that this type of systems is growing at fast pace. In contrast to (Benkhelifa et al., 2014), this work focuses on Vehicular Ad hoc Networks (VANETs), sensors, social networks and Car-to-X, where X can either be infrastructures or other cars. These technologies are shaping the future with the objective of giving a ubiquitous sensing of the surroundings. Later on the work it is identified that these systems produce large quantities of data, changing the context of looking at them from small, simple solving problems, to big data problems that require stronger and more capable algorithms to be solved. Lastly it is presented a problem regarding the interoperability of these systems, which is yet to be solved. The interoperability of these systems is important due to the necessity of presenting a holistic view of the problems in the city.

In the literature there are already some papers that address the need to create a smart emergency system. A good example of this is (Radianti et al., 2014), where the authors present emergency systems and then start to develop a platform that intends to mimic these systems in a smarter way. The authors used a smartphone based publish-subscribe system to accomplish this. The platform helps users by sensing their surroundings and assessing the current disaster scenario, providing them with a safer way to exit the building. It is interesting to analyse the communication that was developed as it takes the data of devices and delivers it, via a web-based broker, to managers and interested parties. The broker also forwards the data to a big database where it is processed in order to retrieve sensor information in useful ways (i.e. charts, reports).

There is also another important topic to cover that is emergency management. According to (Feng and Lee, 2010) it's a process that continuously prepares for disaster even before it happens. It intends to protect people from natural or man-made disasters. It is expected that it can integrate many emergency sources to provide the best possible outcome for the situation. The main purpose of this paper is to explore the possibility of a service-oriented architecture for emergency systems. The authors propose an architecture for this scenario and conclude that these type of systems are of extreme importance in the nowadays world.

In our work we intend to present an architecture for a generic smart system that collects, processes and delivers useful data to users. In the future a smart emergency system will be developed and will integrate information from many places, process it and then retrieve it to interested parties. It is important to understand that this work is a necessary step to accomplish a system with the minimum possible flaws. Also we will integrate technologies that lead us to a more prepared system.

3 SMART CITIES TECHNOLOGIES

Systems related with smart cities require different technologies in order to be fully addressed, therefore this section aims to cover and introduce some of them.

It is important to understand that these types of technologies are of extreme importance in this topic, some of them are directly related to the data collection and storing, while others focus on the processing part of the data lifecycle. Although this section will cover most of them, it will provide more information regarding the processing part.

To begin with, the concept of Big Data (Friess and Vermesan, 2013) shall be addressed. It is understandable that having so many information inputs (sensors, smartphones, etc.) leads to a huge amount of information that needs a new type of treatment.

In (Friess and Vermesan, 2013) the authors refer to big data as "(…) the processing and analysis of large data repositories, so disproportionately large that is impossible to treat them with the conventional tools of analytical databases." The authors also explain that this data is produced by machines, that are much faster than human beings, and according to Moore's Law this data will grow exponentially. Furthermore the authors start pointing out the major contributors for data production (i.e. web logs, RFID, sensor networks, social data, etc…).

It is also referred that Big Data requires different technologies to process the massive amounts of data within a comprehensive amount of time thus, some tools are presented in order to show the current standards in this field.

Additionally, regarding this topic, the authors explain that major companies in the big data topic have a tendency to use Hadoop (Gu and Li, 2013) due to its reliability, scalability and distributed computing.

In (Jara et al., 2014) the authors present a challenge to Big Data, which is of great relevance for our work. This challenge is, perhaps, one of the most important concepts correlated with Big Data not only because of the large amount of data but also because of the IoT paradigm.

The challenge presented is the new way of interaction between humans and the Internet via smart devices. This challenge exists, because of the way that the Internet was created, until now the Internet was based on a human to human kind of interaction, because it delivers content produced by humans for other humans. This kind of communication will not disappear, however new types of interactions will appear as smart objects integrate the nowadays world.

These new types of interactions produce large amounts of data, this is where Big Data comes into play. As has been described in this section Big Data helps us to store this large amounts of data, with the objective of being analysed by intelligent algorithms and tools to extract information and provide knowledge that will empower the applications made recurring to it.

At this point it's possible to conclude that Big Data requires special treatment as it is bigger and contains more information than typical data. For that some algorithms and tools shall be addressed with the intent to choose the most suitable to the presented system.

As posted above major companies around the world to process big data are utilizing Hadoop. Hadoop is a framework that processes big data in a distributed environment (Apache Hadoop, 2014).

Also, it is planned to scale up from single to multiple machines, where each of them provides space and computational power. This framework can also handle failures in applications. It seems like a good way to implement the system. However, in more recent works, despite being around since 2009, Spark (Gu and Li, 2013) started to be used instead of Hadoop.

In (Gu and Li, 2013) the authors made a comparison between the Spark and Hadoop aiming to show which was more suitable for production. It is important to understand that Hadoop is an implementation of the MapReduce framework developed by Google. According to (Gu and Li, 2013) this framework is not designed to support applications with iterative nature, as it cannot keep data during execution time. Because of this, at each iteration, it needs to access disk. On the other hand, Spark, despite being a MapReduce-like framework, is designed to address its current shortcomings regarding iterative applications.

Finally the authors concluded that both frameworks are good, but their application requires a good analysis of the situation. If there is a lot of memory to run the application Spark is definitely faster than Hadoop, on the other hand Hadoop uses less memory but a lot more space in disk.

Other types of data processing are also interesting in the Internet of Things (IoT) context, due to their ability of processing data streams. For instance we can point out Complex Event Processing (CEP) (Chen et al., 2014) and Storm (Toshniwal et al., 2014). Notice that CEP is only a method of analysing and processing streams of data, on the other hand Storm is a distributed computation framework that helps with the processing of large streams of data.

CEP is defined in (Chen et al., 2014) as an effective mechanism that analyses data includes it in a context and triggers events. CEP can, for instance, analyse streams of temperature and determine if the changes in that temperature are normal or abnormal. It can also relate different types of event that lead to a single complex event, such as: (1) flames; (2) temperature spike; (3) sudden humidity decrease. From these three events the system could infer that a fire was happening. Additionally (Chen et al., 2014) aims to develop an architecture for the IoT based on distributed complex event processing. The intent behind distributed CEP is to shorten the bandwidth and the necessary computation.

Storm (Toshniwal et al., 2014) is a real-time distributed stream data processing engine that manages data streams. It was designed to be scalable, resilient, extensible, efficient and easy to administer which makes it a very robust and usable structure. Figure 1 presents a storm topology, which is the real time component that runs all the logic. Topologies are then divided in spouts and bolts. Spouts, represented by the water tap in Figure 1, and are the source of the streams of data. Bolts, represented by bolts on the topology, intend to consume the data sent by spouts, process it and then produces processed outputs.

Furthermore

Figure 1 provides a fault tolerant and scalable architecture for handling data. Additionally this architecture provides the concept of worker that can be interpreted as a node which is programmed to execute a specific task. These tasks may vary, although a good example can be using a worker to process the stream with the Esper queries. In other words each Bolt is associated with a query to be applied in the stream. This will create an efficient a quick way to process the incoming stream and query it for different types of alarming events.

Additionally this two technologies together help one another, in other words Esper needs something to organize and provide data which means that some system needs to be implemented to provide Esper with the data. This is where Storm is useful, it can handle the data management and Esper will handle the queries. This approach will join both systems to enhance both of their main capabilities when dealing with these type of data.

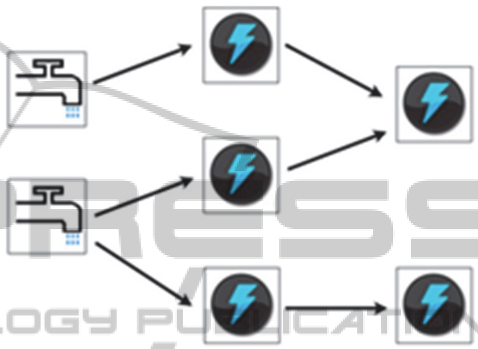

Figure 1: Storm Topology (Apache Storm, 2014).

A very interesting aspect of Storm and CEP, is that they both can work together to provide an excellent way of processing and analysing data in our scenario.

To access this data, sensors and other devices are required. With the intent of making a more transparent communication, the concept of Machine to Machine (M2M) (Wan et al., 2012) emerged. According to (Wan et al., 2012) M2M refers to the automatic communication between, computers, sensors and other devices in the surroundings. This topic is relevant because it makes sensor-to-server communication and sensor-to-sensor possible. This allows the system to constantly check for new data and vice versa.

This concept leads us to another one related to the communication that is publish-subscribe services. According to (Ordille et al., 2009) these services broadcast information to the subscribed parties. In these types of systems a subscriber is a device that will receive information from the publisher. This translates into a much more transparent system, because the publisher can send information to the subscribers and vice versa. Finally in (Radianti et al., 2014) their publishers are treated as the ones that generate information in the form of events. Subscribers are treated as the ones that subscribe to arbitrary flows of information. And brokers are a middle layer between the two participants to pass along the information.

In short, these technologies, due to their relevance in this topic, seem to be an absolute need. They provide a coherent and robust ecosystem to help developers create and deploy their applications. The combination between Storm and Esper seems to be very interesting, since it provides an elegant approach to the topic.

In the latter sections some of these technologies will be addressed again, from an implementation point of view, the main goal is to provide ideas for a future implementation, leaving comments on which technology is the most suitable choice for a specific component of the architecture.

4 USE CASES

In this section current use cases of similar systems will be addressed. This will result in a better knowledge base for the current standards in the area. For this, not only examples of smart cities will be presented but also examples of emergency systems that became smarter with the inclusion of these new concepts.

Lately many smart cities have emerged, such as Amsterdam (Amsterdam Smart City, 2014), Santander (Santander Facility, 2014), Barcelona (Barcelona Open Cities Challenge, 2014), and many others. These cities, due to constant innovation projects and investments, have a tendency to be pioneers in the adoption of new standards in this field. These cities use smart systems help the decision and facilitate the decision making process.

In Finland, the city of Helsinki is running a cooperation cluster called Forum Virium Helsinki (Forum Virium Helsinki, 2014) to provide a platform to develop ICT-based services in cooperation with enterprises, public authorities and citizens as end-users. The platform is concentrated on five project areas, one of them being a smart city initiative focusing on the development of mobile phone services to facilitate urban travelling and living. It also opens up public data so that companies and citizens can create new services by combining and processing the data in innovative ways. This resembles the LivingLab movement that has spread across Europe in the 2000s (The European Network of Living Labs, 2014).

The city of Santander, for instance, uses sensors to monitor the environment, parking areas, parks, gardens and irrigation systems. These sensors are scattered around the city in order to produce alerts that will notify end users with useful knowledge of the situation.

The data is captured by an IoT node that monitors indicators such as temperature, noise or light. This data then travels through repeaters positioned in higher grounds, which send it to the gateways. Lastly this data is stored in a database or sent to other machines where it’s needed.

Regarding the environmental scenario, from a user’s point of view, the available indicators are the temperature, CO level, luminosity and noise, this allows them to receive useful inputs for their wellbeing throughout the day.

The environmental monitoring system is important because it shows how sensors interact with the server and how the server communicates back to the sensors and other subscribers that need this type of information. To summarise, we will discuss the “Participatory Sensing” concept (Description of implemented IoT services, 2014) to obtain a better knowledge about how users interact with the platform, and in which way is it relevant to their day-to-day life. Figure 2 illustrates the concept of participatory sensing from a user’s point of view, which helps us understand how a typical user interacts with this kind of technologies and also how they provide useful inputs to understand the type of data a user needs during application usage. It is possible to visualise that a user can, in this case, publish events, search for events, visualise historical data, subscribe and unsubscribe to events and receive notifications.

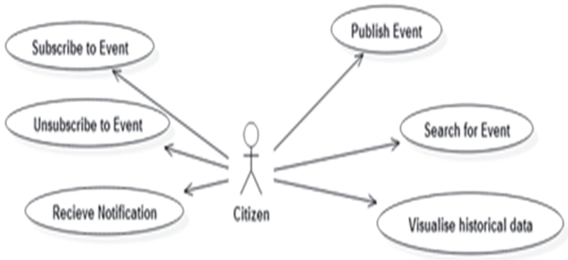

Figure 2: Participatory Sensing - Use Case Diagram [Adapted from (Description of implemented IoT services, 2014)].

The components of the participatory sensing system are: a mobile client for end users to utilise; a server, capable of iterating through data and providing links between the apps and the SmartSantander platform also known as “Pace of The City Server”; and a module that allows devices to register onto the platform. Also, there is a system called “Universal Alert System” (UAS) system, which aims to fire user’s notifications.

The “Participatory Sensing” concept allows users to actively participate in the city ecosystem. The information is then sent to the SmartSantander platform. The concept starts to get even more interesting when users become subscribers of the city systems and are able receive updates of the current status of the city or the road they have to cross to reach their destination. This type of instant real time information directly affects the city from a user’s point of view due to its constant availability and usefulness. The system is available for smartphone via the app and for none smartphone users, via SMS or call.

Additionally Santander city provides other interesting case studies, which are “Precision Irrigation” and “Smart Metering”.

Precision irrigation is a service that intends to provide a useful way of monitoring plants necessities and guarantee that they are fulfilled. Rather than being applied to a whole park, this system is applied by sections or individual plants. Also, the system not only focus on water management but also in other plant needs and their species and growth patterns to minimize the effort from the staff. Even though it looks a bit off the topic this system allowed to realise the necessity of designing the system to accept communications with REST and WebSockets, which are the communication technologies used by it.

Smart Metering system aims to provide IoT based solutions to monitor energy usage in offices. To address this problem new components have been added to the architecture to generate, collect and store the data and information. In addition to these, intelligent components have also been created in order to provide useful information in user-friendly way. These components provide real time analysis of data and consequent knowledge extraction. With this it can identify energy failures and reports on energy consumption that can be drilled down to a specific case.

The last system analysed was (Cecchinel et al., 2014), which is a prototype, named SMARTCAMPUS that aims to equip the SophiaTech campus with sensors to inspire the creation of new applications. Once more the system was chosen due to its usefulness and value in terms of possible inputs for our system.

The SMARTCAMPUS deals with many types of sensors to collect the data. To tackle this challenge the authors propose the architecture seen on Figure 3. This architecture divides in two main focal points: the message collector which intends to collect all data from the internet or sensor networks, to further store in a database that acts as a message queue; and the message processing that aims to process the messages stored in the queue. These components then store the processed information in a database.

Furthermore the architecture contains a configurator, which acts as a routine that can be called periodically to propagate a specific sensor configuration through the network. It also contains a database that contains the current sensor parameters, an API to provide an administrator interface to connect with sensors and a data API that directly accesses data to provide statistics or other types of knowledge.

5 PROPOSED ARCHITECTURE

This section aims to present our architecture to address the typical Smart City scenario. This architecture will provide a way to gather information

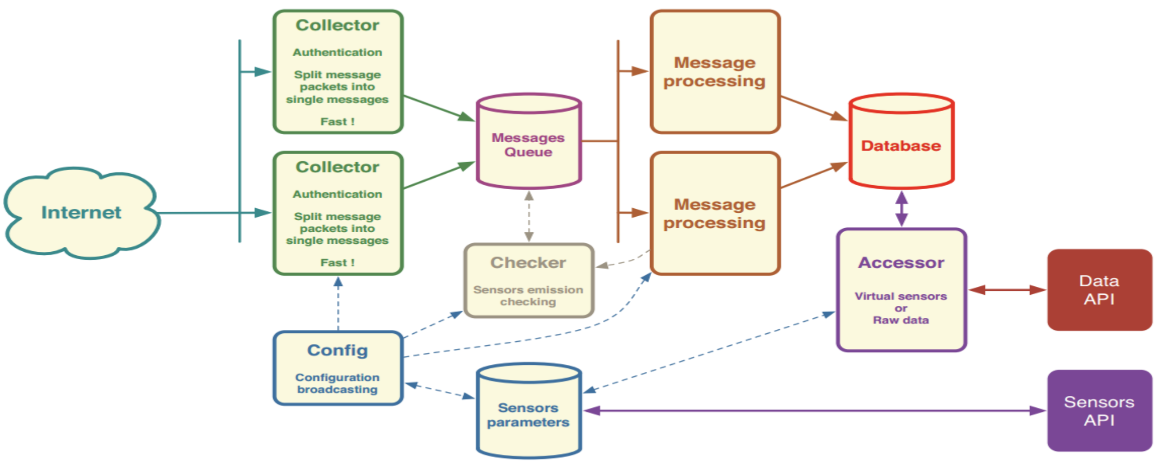

Figure 3: Middleware Architecture (Cecchinel et al., 2014).

from many sources process it and provide useful information to the interested parties.

One of the most important things to understand is that nowadays data comes mostly in streams, which presents an issue due to the tools needed to process it. The tool that we projected to use, to process streams of data is Storm, which has already been documented in this paper. Even though Strom, by itself, cannot retrieve results one hundred percent accurate, due to being stream oriented, we plan to overcome this problem by implementing a parallel processing block with Hadoop. This will, not only provide exact results when the large amount of data is processed, but also provide a better knowledge of the data.

The approach was inspired by the lambda architecture (Lambda Architecture, 2014) with a concrete direction of using the publish/subscribe pattern. The background from other related projects allowed us to perceive that some technologies may not suit very well the collection and direct processing of data. Thus, we opted by a more complex approach that allows to a more scalable and reliable system.

This type of approach also led us to extend the capability of receiving data from multiple sources, which is extremely important in the context of IoT. Furthermore, we shall analyse the proposed architecture, present in Figure 4.

Our architecture is projected to act as an API to provide a connection between data in the IoT and the final user, with the intent of providing relevant information regarding emergency situations.

The system will receive a data stream from IoT nodes, which is then duplicated to be processed by the batch and the speed layer. After that the data is merged with the intent of providing the result with the biggest confidence level associated. When the data is merged a bottleneck can happen, although this situation will be prevented by accepting the first result to appear with the highest confidence level. This can happen in two ways: (1) the stream layer finishes and the batch layer continues to process. With this scenario the stream layer result will be returned with a confidence level attached to it; (2) the stream and batch layer finish at the same time. In this case the data will be merged to provide the most accurate output.

After the data is merged it reaches another processing block, which intends to filter and redirect the acquired knowledge to the subscribed parties. Additionally this block sends the processed data to the statistical data block. The latter block not only keeps track of statistical data to help us understand patterns along the year but also provides data to construct KPI's, charts and reports.

After the processing is all done, users can access the data in two ways: (1) via the data API, which is projected for developers who want to build applications around this context; (2) via the data output, which will serve to return the data to the subscribed parties. Additionally the API will provide a way of notifying other sensors in the field, which means that if a sensor sends a fire alert, other sensors around it will be asked for their current situation to localize the hazard with maximum precision. This type of communication is also important if the fire is located near a road since the system can be prepared to notify street lights to prevent drivers from entering the affected road. Also in the highways a lane can be closed and the traffic redirected to other lanes or even roads.

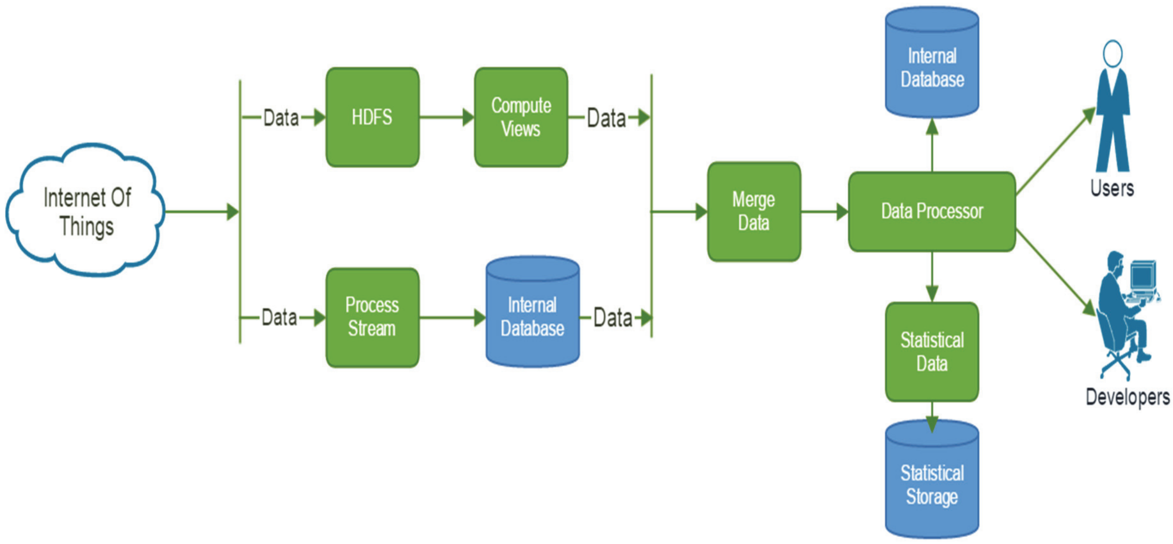

Figure 4: Proposed architecture.

This architecture can be applied in many different scenarios; one of them will be addressed so that we can establish an example to explain some of its functions. Let's assume we have three types of sensors: smoke, flames and temperature. These sensors are constantly sending a stream of data into our system, the idea is to process this data in order to figure out whether we are in the presence of a fire or not. The system has a threshold that serves as a maximum possible value for a normal event, when crossed they trigger events that can lead to, in this case, a fire. Having different types of sensors allows us to better understand whether the fire is happening. Different combinations of events can occur, thus the system must have something to divide the ones that are indeed problematic. Furthermore we shall materialise this example:

- If there is smoke, flames and the temperature passes the threshold, then we have a fire;
- If there is smoke, no flames and the temperature is rising, it is possible to have a fire.

Many more combinations can be presented, although these explain the concept that we are trying to achieve.

Furthermore each module of the presented architecture should be accounted for when choosing the right technologies, in order to access the full potential of it. Hence we need to account for the data stream that is arriving. For instance, it should use a publish-subscribe messaging system, which will handle the stream and split it into events that can be processed by the rest of the modules. The events that have been split will be processed by the both layers. At this point, in the speed layer, there are two important things to acknowledge: (1) it is advised to use a complex event processing system due to the nature of the system, this will provide an event based approach which will necessarily climax with event correlations and a smarter way of dealing with the data stream that constantly change. This approach will also provide the ability of integrating many types of events at once, this will expand system acceptance in terms of receiving events and inevitably prepare it to explore further sensor integrations; (2) an in memory database for storing alarming events is also useful, because of the high demand from the system.

In the batch layer algorithms with predictive capabilities should be added to enhance the system overall quality and usefulness. This will provide ways to calculate KPI's, draw charts and predict whether it is important or not to be in maximum alert level. From a high level perspective this type of inputs seem to have a great importance, with applications such as divide a specific fire protection team to a zone which is prone to peaks of fire during the summer or redirect traffic because a particular road is more likely to be affected by the floods in the winter.

The rest of the processing components in the system can be executed with any programming language and should withstand the volume and velocity of data, also the code statements should be optimized to minimize overheads and bottlenecks. The databases should be chosen according to the needs of each specific scenario. It is important to understand that many database systems can be chosen to incorporate the solution, although for each specific situation a brief analysis of the problem should be made in order to perceive the best possible choice. As a practical example we can point out that the database in the speed layer should be in-memory, on the other hand the statistical storage could be an NO-SQL database that supports large quantities of data to enhance overall system scalability and has a good read mechanism due that its main focus is reads.

Moreover other important aspect to discuss is the communication. The way the system is designed, and from the lessons learned from the use cases, the best technologies should be REST, WebSockets and MQTT. REST will provide an easy and consistent way to access the API, providing endpoints for events and the ability to execute filters in the queries. WebSockets are useful because due to the facilitations in terms of real-time communication. The MQTT protocol is important to establish connection between the system and sensors and actuators scattered in the city in order to extract real-time information.

Additionally other important aspect is the inclusion of a message broker, which will accepts messages from the source divide the stream of data in messages that are easier to process and correlate for a better, more useful and more accurate output, which is delivered to a consumer.

6 CONCLUSIONS AND FUTURE WORK

This paper intends to document the current state of the art in smart city systems and their related technologies. Over the coming sections the problem has been documented and some use cases were studied to provide the most possible inputs with the intent to understand which challenges existed and needed to be tackled.

The analysed documents provided several useful outputs to establish a good baseline in terms of architecture and tools to be used. For instance the concept of participatory sensing, from SmartSantander, led us to think that with so much user data available the system could be adapted to process it with the intent of retrieving knowledge from it. Another example of a good tool to process data in IoT is the lambda architecture, which provides the best of the stream layer processing allied with the batch layer that provides more accurate results.

The knowledge extracted from the state of the art systems and technologies guarantees that our contributions were, as expected, scalable, adaptable, feasible and viable.

Furthermore, we aim to develop a system that will address the current shortcomings in this context. This system will be more directly related to emergency management. Therefore we aim to construct a platform that receives disaster data from many sources, process it via established components and lastly retrieves it to any party that subscribed to the specific type of event. Consequently this paper also serves as a document to establish an architecture for that type of system, serving as a first practical application of it. An initial overview of the technologies that can be used was also made with the intent of providing the necessary steps to implement a similar system, or at least provide some additional knowledge regarding this topic.

A smart emergency system is important in the current context due to its usefulness and transparency while dealing with data, as it can provide predictions and problems before they happen to managers. Thus, with the use of this type of system data becomes clearer and leads to a more prepared and quicker response to any emergency or disaster.

Another interesting application, which empowers the system, is social mining, which due to the importance of social networking in nowadays society seems like and excellent way to complement the inputs of the system.

This is important to complement the system because it can detect disasters via a post in a social network. The post does not need to be in a specific format, the algorithms will only be looking for keywords that will trigger the attention of the system. Although this data is extremely relevant, it is important to guarantee that it isn't false. A possible solution for this problem can be a request to the sensors that are placed in that specific site.

In short, Internet of Things is successfully thriving in the current world, therefore these type of systems will continue to emerge alongside it. An excellent way to evolve and prepare future cities is to be more interconnected and aware, in essence enabling better decision-making.

ACKNOLEDGEMENTS

This work was partially financed by iCIS – Intelligent Computing in the Internet Services (CENTRO-07-ST24 – FEDER – 002003), Portugal.

This work was also made possible with the help of Ubiwhere, Lda, which provided useful inputs in discussions and also the facilities.

REFERENCES

Alazawi, Z., Alani, O., Abdljabar, M. B., Altowaijri, S., and Mehmood, R.. 2014. A smart disaster management system for future cities. In *Proceedings of the 2014 ACM international workshop on Wireless and mobile technologies for smart cities (WiMobCity '14)*.

Albtoush R., Dobrescu R., Ionescou F., (2011) A Hierarchical Model for Emergency Management Systems.

Amsterdam Smart City, (2014). [online] Available at: http://amsterdamsmartcity.com/?lang=en.

Anttiroiko, A., Valkama, P. and Bailey, S. J., (2014). Smart cities in the new service economy: building platforms for smart services. *AI Soc. 29*(3): 323-334

Apache Hadoop, (2014). [online] Available at: http://hadoop.apache.com.

Apache Storm, (2014). [online] Available at: http://storm.apache.com.

Aurigi, A. (2005) Making the digital city. The early shaping of urban internet space. Ashgate, Aldershot

Barcelona Open Cities Challenge, (2014). [online] Available at: http://opencities.net/barcelona.

Benkhelifa, I., Nouali-Taboudjemat, N. and Moussaoui, S. (2014). Disaster Management Projects Using Wireless Sensor Networks: An Overview. *2014 28th International Conference on Advanced Information Networking and Applications Workshops*.

Carillo FJ (ed) (2006) Knowledge cities. Approaches, experiences, and perspectives. Elsevier, Amsterdam.

Cecchinel, C.; Jimenez, M.; Mosser, S.; Riveill, M., (2014). An Architecture to Support the Collection of Big Data in the Internet of Things Services (SERVICES).

Chen, C., Fu, J., Sung, T., Wang, P., Jou, E. and Feng, M. (2014). Complex event processing for the Internet of Things and its applications. *2014 IEEE International Conference on Automation Science and Engineering (CASE)*.

Chourabi, H., Nam, T., Walker, S., Gil-Garcia, J., Mellouli, S., Nahon, K., Pardo, T. and Scholl, H. (2012). Understanding Smart Cities: An Integrative Framework. *2012 45th Hawaii International Conference on System Sciences*.

Description of implemented IoT services, (2014). [online] Available at: http://smartsantander.eu/downloads/Deliverables/D4.2.pdf.

European Network of Living Labs, (2014). [online] Available at: http://www.openlivinglabs.eu/

Feng, Y. and Lee, C. (2010). Exploring Development of Service-Oriented Architecture for Next Generation Emergency Management System.*2010 IEEE 24th International Conference on Advanced Information Networking and Applications Workshops*.

Forum Virium Helsinki, (2014). [online] Available at: http://forumvirium.fi/en

Friess, P. and Vermesan, O., Internet of Things: Converging Technologies for Smart Environments and Integrated Ecosystems. Aalborg, Denmark: River Publishers, 2013.

Gu, L. and Li, H., "Memory or Time: Performance Evaluation for Iterative Operation on Hadoop and Spark," *High Performance Computing and Communications & 2013 IEEE International Conference on Embedded and Ubiquitous Computing (HPCC_EUC), 2013 IEEE 10th International Conference on*, vol., no., pp.721,727, 13-15 Nov. 2013.

Hollands RG (2008) Will the real smart city please stand up? Intelligent, progressive or entrepreneurial? City 12(3):303–320

Jara, A.J.; Genoud, D.; Bocchi, Y., "Big Data in Smart Cities: From Poisson to Human Dynamics," *Advanced Information Networking and Applications Workshops (WAINA), 2014 28th International Conference on*, vol., no., pp.785,790, 13-16 May 2014 doi: 10.1109/WAINA.2014.165

Komninos N (2002) Intelligent cities. Innovation knowledge systems and digital spaces. Spon Press, London

Lambda Architecture, (2014). [online] Available at: http://lambda-architecture.net/

Ordille, J., Tendick, P. and Yang, Q. (2009). Publish-subscribe services for urgent and emergency response. *Proceedings of the Fourth International ICST Conference on COMmunication System softWAre and middlewaRE - COMSWARE '09*.

Piro G., Cianci I., Grieco L. A., Boggia G., and Camarda, P. 2014. Information centric services in Smart Cities. *J. Syst. Softw.* 88

Radianti, J., Gonzalez, J. and Granmo, O. (2014). Publish-subscribe smartphone sensing platform for the acute phase of a disaster: A framework for emergency management support. *2014 IEEE International Conference on Pervasive Computing and Communication Workshops (PERCOM WORKSHOPS)*.

Santander Facility, (2014). [online] Available at: http://www.smartsantander.eu/index.php/testbeds/item/132-santander-summary

Toshniwal A., Taneja, S., Shukla, A., Ramasamy, K., Patel, J. M., Kulkarni, S., Jackson, J., Gade, K., Fu, M., Donham, J., Bhagat, N., Mittal, S., and Ryaboy, D., Storm@twitter. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data (SIGMOD '14)*.

Vakali, A., Anthopoulos, L. and Krco, S. (2014). Smart Cities Data Streams Integration. *Proceedings of the 4th International Conference on Web Intelligence, Mining and Semantics (WIMS14) - WIMS '14*.

Wan, J., Li, D., Zou, C. and Zhou, K. (2012). M2M Communications for Smart City: An Event-Based Architecture. *2012 IEEE 12th International Conference on Computer and Information Technology*.