Received May 31, 2017, accepted June 25, 2017, date of publication July 5, 2017, date of current version August 14, 2017.

Digital Object Identifier 10.1109/ACCESS.2017.2723659

Practical Lessons From the Deployment and Management of a Smart City Internet-of-Things Infrastructure: The SmartSantander Testbed Case

PABLO SOTRES, JUAN RAMÓN SANTANA, LUIS SÁNCHEZ, JORGE LANZA, AND LUIS MUÑOZ, (Senior Member, IEEE)

Network Planning and Mobile Communications Laboratory, Universidad de Cantabria, 39005 Santander, Spain

Corresponding author: Pablo Sotres (psotres@tlmat.unican.es)

This work was supported in part by the research project SmartSantander through the 7th Framework Programme of the European Commission under Grant FP7-ICT-2009-5 and in part by the Spanish Government by means of the project ADVICE "Dynamic provisioning of connectivity in high density 5G wireless scenarios" under Grant TEC2015-71329-C2-1-R.

ABSTRACT The smart cities vision is inexorably turning into a reality. Among the different approaches used to realize more intelligent and sustainable environments, a common denominator is the role that information and communication technologies will play. Moreover, if there is one of these technologies that emerges among the rest, it is the Internet-of-Things (IoT). The ability to ubiquitously embed sensing and actuating capabilities that this paradigm enables is at the forefront of the technologies driving the urban environments transformation. However, there are very little practical experiences of the IoT infrastructure deployment at a large scale. This paper presents practical solutions to the main challenges faced during the deployment and management of a city-scale IoT infrastructure, which encompasses thousands of sensors and other information sources. The experience we have gained during the deployment and operation of the IoT-based smart city infrastructure carried out at Santander (Spain) has led to a number of practical lessons that are summarized in this paper. Moreover, the challenges and problems examples, excerpted from our own real-life experience, are described as motivators for the adopted solutions.

INDEX TERMS Data quality, deployment, Internet-of-Things, smart city.

I. INTRODUCTION

The Internet-of-Things (IoT) concept has attracted a lot of attention from the research and innovation community for a number of years already [1]–[4]. One of the key drivers for this hype towards the IoT is its believed applicability to a plethora of different application domains [5], like e-health [6], [7], smart-environment [8], smart-home [9] or Industry 4.0 [10]. However, there is one application for IoT that is probably standing out from the rest, Smart Cities [11], [12]. Different city-domain stakeholders (technicians, city planners, politicians, researchers, etc.) will need to implement actions aimed at assuring that some key quality criteria related to the sustainability and efficiency in the city domain are fulfilled [13], [14]. However, although the available related literature is large and encompassing many different approaches [15], [16], there is a particular lack of practical approaches that real practitioners could use to learn from (i.e. "standing on the shoulders of giants" Sir Isaac Newton dixit).

Interestingly, there has been some initiatives that, in order to improve these solutions' maturation and significant roll-out, have tried to support the evaluation of IoT solutions under realistic conditions in real-world experimental deployments [17]. However, still they tend to lack from the necessary scale or they fail to fulfil some key indicators. Conversely, there are little practical references to the main aspects of deploying and managing (even fewer in this latter aspect) a city-scale IoT infrastructure. Existing works focus on the high-level view and applications perspective [18]–[21], but they do not delve any deeper on infrastructure deployment and management issues. Thus, there are plenty of significant technological advances and also several real-world deployments, but there is, to the best of our knowledge a lack into bringing the two approaches together and present the practical challenges posed by the deployment and management of a Smart City IoT infrastructure and the technological solutions addressing that challenges in a way that practical lessons are described so practitioners can stand on them.

2169-3536 © 2017 IEEE. Translations and content mining are permitted for academic research only.
Personal use is also permitted, but republication/redistribution requires IEEE permission.
See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

The main contribution of this paper is presenting, from a practical standpoint, the lessons learnt during the deployment and operation of the SmartSantander testbed [22] bringing together challenges and solutions. The SmartSantander infrastructure includes a continuously growing IoT setup spread throughout the city that currently encompasses more than 10,000 diverse IoT devices (fixed and mobile sensor nodes, Near-Field-Communication (NFC) tags, gateway devices, citizens' smartphones, etc.). Fig. 1 shows some examples of these deployed devices.

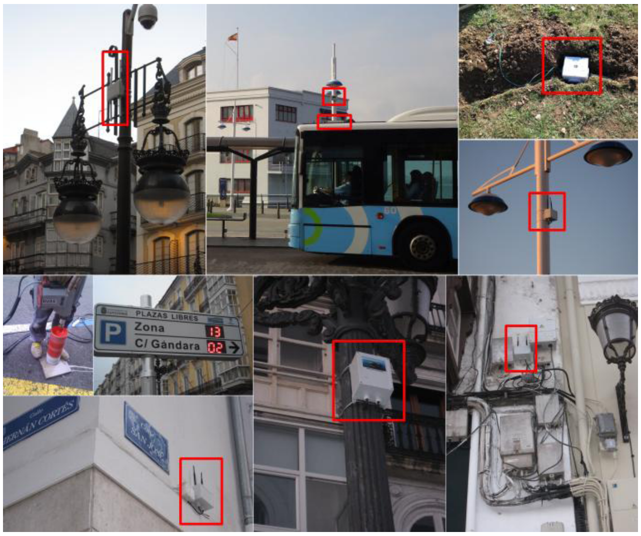

FIGURE 1. Real SmartSantander deployment examples.

In this sense, the paper focuses on two key aspects that Smart City technical managers will face during the lifecycle of their IoT infrastructure set-up and operation. Firstly, the practicalities of the physical deployment of the Wireless Sensor Network (WSN) that conforms the core of the SmartSantander infrastructure and the extension of IoT infrastructures beyond the use of that WSN. Including the need for interoperability in a highly heterogeneous system. Secondly, the capacity to continuously monitor the health of the IoT infrastructure both in terms of the underlying devices and in terms of the data that they are constantly generating.

The remaining of the paper is structured as follows. In Section II we present a non-exhaustive overview of the existing related work around the management of large-scale IoT infrastructures and the technologies that comprises the baseline of the solutions that we have adopted for the monitoring of the SmartSantander infrastructure. Section III describes briefly the SmartSantander infrastructure. In this sense, several articles already provide details about the composition of the deployment made in Santander so Section III will provide just an overview and the relevant references for the sake of self-completeness of this paper. The practical lessons extracted during the deployment and operation of the SmartSantander infrastructure will be presented and described in Section IV. Section V delves into the practical approach adopted for infrastructure management and monitoring. Finally, Section VI concludes the paper highlighting the key contributions and summarizing most relevant practical findings resulting from the practical work that we have been developing for the last five years.

II. RELATED WORK

The successful deployment of a wireless sensor network is a difficult task, littered with traps and pitfalls. Even a functional network does not guarantee gathering meaningful data. As it has been already introduced the two aspects that will concentrate the focus of this paper relates to the deployment of the SmartSantander IoT infrastructure and to the monitoring of its status and quality of the provided data. In this section we make a non-exhaustive review of existing literature around these two topics and how the contributions presented in this paper differentiates from them.

A. IoT INFRASTRUCTURE DEPLOYMENT

In [17] a summary of experimental IoT facilities was presented and several key aspects were analyzed for each of them. Among these requirements, scale, federation and end-user involvement are particularly relevant when the objective of an IoT infrastructure exceeds the experimentation scope and they are meant to support real service provision.

Some large scale IoT testbeds have been deployed lately. However, they lay on the comfortable indoor environment [23] or even they are generated virtually [24]. While some practical deployment aspects can be learnt and extrapolated to a real IoT deployment, their value in this respect is limited. The scale and reality of the SmartSantander deployment has enabled us to derive some interesting practicalities that can be used by IoT practitioners on their real-life developments.

Since IoT is a quite novel paradigm, many different competing technologies have been struggling to take lead. Hence, one of the known issues for closing the gap between innovation and big-market penetration [25], is the interoperability of the already different alternatives. Especially important, for achieving real interoperability, is the information modelling when the IoT infrastructure is to be formed by a heterogeneous combination of systems, which is the most probable situation in future real scenarios. Semantic technologies [26] are being explored to provide a common ground in this modelling effort. However, proposed solutions [27]–[29] have been only been proposed from a theoretical standpoint and only very recently application of semantics to real deployments have been explored in a proof-of-concept approach [30]. The solution adopted in the SmartSantander deployment sets some baseline for the use of formal semantics as it has defined specific models for the smart-things and the information they provide, but it was decided to not use semantic ontologies due to the additional computational burden that they imply.

Often smart city projects have a top-down approach focused on improving city infrastructure using technology for specific application domains. However, there are evidences that grass-roots based Smart City projects deliver better value and sustainable success [31]. In this sense, SmartSantander platform is delivering the enablers to promote this participation and federate IoT infrastructure that might come from large utilities [32] to individuals [33].

B. IoT INFRASTRUCTURE HEALTH MONITORING

In the IoT, information obtained from a large-scale infrastructure of smart-things, is where intelligent decision-making processes and added-value services are rooted. If devices producing this data are not reachable or observations gathered contain poor-quality data, services and decisions are likely to be flawed. IoT infrastructure health monitoring is critical to achieve real engagement and acceptance of the IoT paradigm and services and transform the current hype into stable and profitable market [34].

In [35] authors survey over the most well-established categorization of Data Quality (DQ) and describes eight dimensions for assessing the quality of data streaming environments namely accuracy, confidence, completeness, data volume, timeliness, ease of access, access security and interpretability. We have used these categories in order to assess the health of the SmartSantander IoT infrastructure. The ease of access, access security and interpretability are non-functional requirements that cannot be monitored objectively. However, as previously mentioned interpretability has been taken into consideration when defining the information models. Moreover, data volume dimension is transversal to the remaining four categories (i.e. accuracy, confidence, completeness, and timeliness), that are the ones on which SmartSantander infrastructure health monitoring mechanisms mainly focuses. In this respect, the two techniques that provides the most meaningful information are outlier analysis [36] and data cleaning [37].

The first one, outlier analysis, has been widely researched in various disciplines [38] and it has the capacity to serve both application domain features (e.g. event or intrusion detection) as well as an important aspect of the IoT infrastructure management duties, namely fault detection [39]. From the myriad of different approaches already available in the literature [36], we decided for a practical, yet simple approach that will be defined in Section V. The main considerations used when implementing the fault detection mechanism for SmartSantander infrastructure were: i) to avoid additional in-network computation in order to minimize (almost eliminate) the burden that sensors have to take due to management duties. The idea is to include in the sensors duty-cycle tasks that only relate to the service they are providing; ii) to define a convenient neighborhood range and a temporal window since we were using spatial and temporal correlation; iii) to consider the mobility of some of the sensors to optimize the outlier detection and to tailor the mechanism to mobile nodes own idiosyncrasy; and iv) to define an appropriate dynamic threshold to best determine outliers.

The second one, data cleaning, is part of the data's life cycle [40] and goes a step further, taking action and removing the anomalous data from the stream. Following a similar approach to the one described in [41], SmartSantander data cleaning techniques are implemented as part of the platform middleware so that application developers can transparently interact with an already cleaned dataset. As it happens with outlier detection, this is a field that have been explored extensively. Similarly to the previous case, we followed analogous practical considerations for implementing the data cleansing modules in the SmartSantander platform. This way, we, again, avoided in-network solutions [42], [43] that put on the actual devices producing the data any of the responsibilities of data analysis. In this sense, a context aware model-based technique [44] has been implemented within the cloud-based SmartSantander platform. In some of the cases, the model is substituted by well-known and reliable sensors which are used as reference [45].

III. SMARTSANTANDER IoT INFRASTRUCTURE

The SmartSantander project [46] targeted the creation of a European experimental test facility for research and experimentation on architectures, key enabling technologies, services and applications for the Internet of Things (IoT) in the context of a smart city.

However, this testbed goes beyond the experimental validation of novel IoT technologies. It also aims at supporting the assessment of the socio-economical acceptance of new IoT solutions and the quantification of service usability and performance with end users in the loop [47].

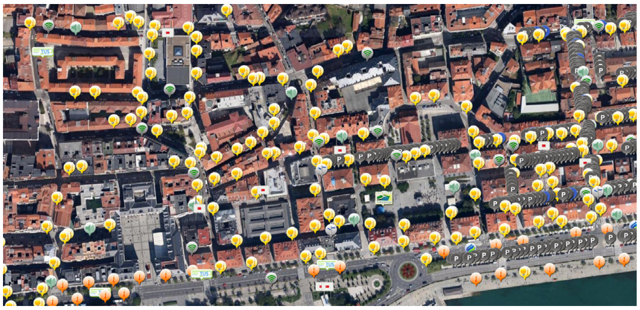

FIGURE 2. Santander IoT infrastructure deployment excerpt view.

Fig. 2 shows an excerpt view of the deployment. The different markers represent the deployed nodes (e.g. illuminance, sound pressure level, ambient temperature, mobile nodes or car presence detection sensors).

The IoT experimentation facility deployed in Santander was selected using a cyclic approach. The deployment, influenced by Santander Municipality's strategic smart-city service requirements, intentionally provided a concentration of IoT devices in the city center (a 1 Km^2 area) in order to achieve the maximum possible impact on the citizens. Nonetheless, other city areas are also covered.

The SmartSantander IoT infrastructure has a core WSN deployment centered on the environmental monitoring application domain [22]. Part of these WSN is installed on vehicles that move around the city, so mobility of devices is at the same time an opportunity and a challenge [48]. Another part of the installed WSN is focused on traffic (one of the main headaches for any city) that imposes tough conditions in terms of radio propagation as devices have to be buried under the asphalt [49].

Nevertheless, IoT goes beyond WSN and, as such, the SmartSantander infrastructure encompasses other sources of information outside the WSN realm. Firstly, a large number of tags (NFC-based) have been deployed and linked to an Augmented Reality application. However, it is precisely the smartphones that, almost, everybody has in their pockets, and some Apps that have been developed, the part of the SmartSantander setup that lends a distinguishing feature to the whole infrastructure. Through these Apps the citizens (through their smartphones) become sensing devices, not only by allowing access to their smartphones' embedded sensors (e.g. compass, microphone, luminosity) but also by reporting events they have observed (e.g. broken bench, malfunctioning streetlight, etc.).

Finally, it is important to mention that cities already have a large amount of data that, in most of the cases is not publicly available due to lack of infrastructure rather than because of privacy/confidentiality policies. Moreover, in the cases that these datasets are opened, there is a lack of the necessary interfaces to actually extract the value from that information [50]. This was the case of Santander.[1] The SmartSantander platform brought a number of public, or publishable, datasets into a common IoT-as-a-Service system. This way, also the legacy devices used to gather all this information are considered part of the infrastructure to be managed.

IV. SMART CITY INFRASTRUCTURE DEPLOYMENT IN PRACTICE

As a result of the experience that we have gathered during the different deployment phases we have undertaken throughout the city of Santander over the last years, we have found some practical aspects that should be considered by anyone targeting a real-world large-scale smart city deployment. These practicalities, not only related to the traditional laboratory-versus-real-world behavior deviations, eventually led to an unexpected increment of the deployment tasks involved efforts and budget. While the reader might see some of those lessons learnt as quite obvious ones, recommendations provided by manufacturers and specifications are full of pitfalls for the unwary.

This section presents those experiences categorized on different domains, based on the specific deployment topic they address. Fig. 3 shows a conceptual representation of all of them.

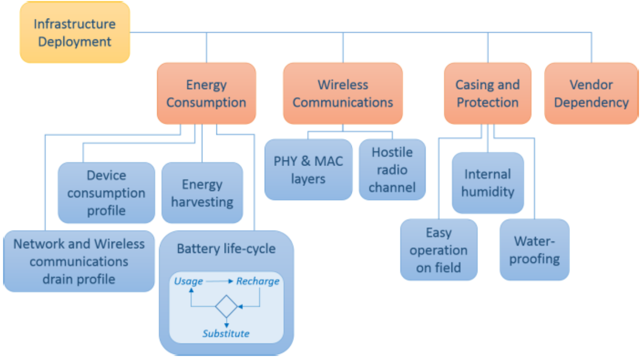

FIGURE 3. Addressed deployment related challenges.

A. ENERGY CONSUMPTION RELATED PRACTICALITIES

Although energy shortage is one of the well-known design principles for IoT, many existing IoT deployments are indoor-based. Thus, since IoT devices can be continuously powered, energy constrains are not usually one of the major factors involved on their design. On the contrary, power consumption and energy harvesting are very important factors to take into account on an outdoor large-scale scenario deployment. In this sense, SmartSantander infrastructure is also heterogeneous in terms of powering mechanisms. IoT devices deployed in the city of Santander can be categorized in:

- Those installed on public buildings without power restriction, hence continuously powered (24/7).
- Those installed on public vehicles, such as taxis or buses, continuously powered by the vehicle battery. Those devices may only work when the vehicle engine is on. In some cases they are continuously powered in order to avoid sensor recalibration problems.
- Those connected to the public lightning network, hence powered when the network is energized. This is the case of most of the devices installed on facades or lampposts. Although this can be different on different deployments, in the case of the city of Santander those devices work using a mixed duty-cycle: during daylight they work on batteries, whereas overnight they are continuously powered by the lightning network and the batteries are recharged.
- Those not connected to any public power source, hence working on batteries using a full duty cycle. This is the case of all the on-street parking sensors, which are buried under the asphalt, and some of the devices deployed on facades and lampposts.
- Those passively powered on demand by external actuators, hence deployed without any kind of energy source. This is the case of all the NFC and QR tags deployed around the city to support the augmented reality and smart shopping use case.

This diversity has allowed us to analyze what are the practical issues imposed by energy constrains on a real-field smart city.

1) ENERGY SOURCES HARVESTING AND THEIR FEATURES

Firstly, the availability of power sources in a city scenario needs to be considered. Of course, street-lighting grid is almost ubiquitous in any city of a developed country, but if

[1] Santander Open Data portal. [Online]. Available: http://datos.santander.es/ Accessed on: April, 26, 2017.

a practitioner wants to start deploying any kind of device on the city, municipality willingness and collaboration is imperative. Still, even with the access to the grid granted, the requirements to isolate their ''critical and proven-to-be-stable'' systems from our ''on-the-edge and not-so-tested'' devices, imposed by the electrical engineers in charge of this infrastructure, can be challenging. In our case, this caused an increase of the cost of each deployed node between 10-25% due to the need to include a power protection switch to every IoT device plugged to the grid. Still this solution was better than using other energy harvesting specific hardware (e.g. solar panel, etc.).

Once the permission to connect our IoT devices to a public grid is achieved, it is still necessary to know about the time behavior and evaluate the different available power sources versus the requirements of deployed nodes. As previously mentioned, not all the accessible power lines will be constantly powered, so energy harvesting might still be needed. If this is the case, then there are two aspects to consider: device power consumption and battery characteristics.

2) WIRELESS TRANSMISSION ENERGY HUNGER

In a sensor-based IoT infrastructure, energy consumption is mostly driven by over-the-air wireless transmission. Therefore, the amount of information that can be sent to the air by battery-powered devices is not only limited by bandwidth, but is also limited by the battery capacity. Several considerations and optimizations can be evaluated depending on the energy duty cycle used by each deployed device.

Among all of them, optimizing when and for how long the device is out of its deep sleep state can significantly contribute to larger battery lifetimes. In fact, this is one of the few enhancements which can be applied independently of the chosen duty cycle. However, still a trade-off on the functionality offered has to be respected. For example, the behavior of our deployed on-street parking infrastructure has gone through three successive iterations to improve the overall behavior (cf. section IV.B) that affected its power consumption. Devices provided by the manufacturer selected during the project (back in 2010) required to be continuously on to detect real-time parking spot status change. As the energy consumption was too high, it was decided to periodically sample the parking spot status and only report a new observation if there was a change. The rest of the time the device was in deep sleep mode to maximize battery duration. This solution just provided near-real-time information, which was evaluated as acceptable for the selected application (parking spot monitoring).

3) BATTERY USAGE AND REPLENISHMENT

On the other hand, when a device needs to be running in a continuous mode but the energy supply is intermittent, the maximum battery charge is limited by the external energy supply active period and by the own battery charging circuit. If the amount of energy charged during the active periods of the external energy supply is not enough to power the device while the external energy supply is inactive then the requirements cannot be fulfilled on that specific deployment point. Furthermore, if the sequential recharge/discharge cycles are not conveniently adjusted, then, after several negative-balance discharge/recharge rounds, the device can reach a ''point of no return'' in which the device is not able to boot or is not remotely accessible. Once a device gets into this situation, the only way of bringing it back to operational mode is to physically recharge and/or replace it on the field. This have a non-negligible cost that affects the maintenance budget. This situation was addressed by introducing a ''low consumption self-healing mode'' threshold in which the device does not generate any new observation until the battery level does not reach a certain value.

Trying to solve these problems by just increasing the battery capacity may be helpful in the future (once the battery maximum capacity degrades). However, it does not have any influence in the devices' stationary behavior, since it would only imply a higher number of negative-balance rounds until the devices get inoperative. On the other hand, increasing the battery capacity usually means a size increment, which also has impact on the device enclosure design and its weight. Moreover, capacity is just one of the important characteristics of the batteries, but charging circuit maximum energy flow is as important and is normally not sufficiently observed. It is also worth mentioning that daylight duty-cycles change depending on the season, thus the minimum active duty-cycle must be considered while designing the battery capacity requirements.

4) DEVICES NETWORKING ENERGY CONSIDERATIONS

Last but not least, network topology can also have an impact on the energy consumption. Even if nowadays there are different emerging Low Power Wide Area Network (LPWAN) technologies, such as LoRa or NB-IoT that matches most of the requirements for outdoor environments, when SmartSantander deployment started back in 2011, the predominant solutions were based on the IEEE 802.15.4 standard. Due to this, the network deployed on Santander is using a multihop mesh-network approach. Still, IEEE 802.15.4 based technologies are valid and being used in more recent deployments. A multihop network is the only solution to achieve large coverage areas out of Wireless Personal Area Network (WPAN) technologies, but it implies the need for intermediate nodes to be on operational state in order to act as repeaters. In our case, we decided to configure those nodes working only on batteries (with no access to energy harvesting solution) as final nodes. This way, all these devices do not behave as repeaters and their battery lifetime duration is extended. However, it is of critical importance to have a sufficiently redundant network because the lack of a valid route to a sink can increase the energy consumption of the isolated devices, as the number of transmission retries will also be increased.

5) VISITING BATTERY-EXHAUSTED NODES

As a final remark of this subsection, civil work required to replace empty batteries in a real scenario need to be scheduled in advance and should not be underestimated. Our experience shows that battery durations on technical specs are often too optimistic, so it is better to have contingence plans in case something unexpected arises and don't try to handle each deployed device as a single one, but focus on a set of devices.

B. WIRELESS TRANSMISSION RELATED PRACTICALITIES

Capacity to wirelessly transmit their observations is the critical functional feature that IoT sensor devices must have in smart city deployments. The key conclusion that we can conclude from the experience we had is that considering local conditions and carefully investigating how local environmental conditions will affect your deployments is a must when planning the deployment of an IoT infrastructure. Real conditions differ from laboratory ones as there are many elements that can be hardly simulated or emulated. Indeed, even though vendors perform their tests under "realistic" conditions, they just mean "real life controlled conditions". Examples as the ones described below, are there to remind us that vendors cannot test their products to check every requirement we might have.

1) PHY AND MAC LAYER ISSUES

As stated before, the network deployed in the city of Santander was based on the IEEE 802.15.4 working on unlicensed frequency band. In Europe this means both 868MHz and 2.4GHz bands can be used. Looking for maximizing the available bandwidth (following the manufacturer's advice), the 2.4 GHz option was used. As mentioned before, we configured the network with high enough device density and Line-of-Sight (LoS) communications between adjacent nodes were almost granted. We also took the precaution to select different channels for adjacent clusters so that inter-cluster interference is minimized at the boundaries of each cluster.

However, even though we followed aforementioned well-known wireless network deployment best-practices, there were still some hidden pitfalls worth mentioning. One of the lessons we learnt arose with the on-street parking infrastructure. Even though the manufacturer stated that the devices have been tested under real conditions, the behavior in our own tests revealed large Packet Error Rates (PER) when sending their observations to the corresponding repeater whenever a vehicle was parked on the parking slot they were monitoring. Some vehicle models produced larger power loss and the node was not able to reach its closest repeater. In addition to the loss in functionality (observations from those nodes were not available), this has another negative side-effect. IEEE 802.15.4 MAC protocol defines an explicitly-acknowledged transmission mechanism with retransmissions upon unacknowledged transmitted frames. In the previously described situation the number of transmission retries the node has to do makes the battery to be drained at a higher rate. In an attempt to save battery we limited the number of retransmissions. However, this only partially patched the consumption side-effect. Still, we had to force periodic transmission of parking lot occupancy information, even if the state was the same (hence slightly increasing power consumption). This way, it was possible to infer the presence or absence of a vehicle even if no observation was received, because this was most likely produced by the presence of a vehicle parked on top of the sensor. We also trialed with some variations just forcing the status notifications every n periods in order to save as much battery as possible, although the result was not yet as accurate as desired. In the long run, we decided to start with the deployment of a new generation of devices using 868 MHz band instead of replacing the batteries of the already deployed devices once they were exhausted.

In order to compensate the attenuation introduced by vehicles parked on top of the nodes, we used the 868MHz frequency band, which has better power loss figures, and created a new parallel cluster to handle these devices. In addition, and thanks to development of sensors' technology during the last years, the communication is event-based and the device only notifies when the parking spot status changes. The result was a much more time-responsive parking monitoring infrastructure that did not suffer transmission losses due to parked vehicles.

An interim solution that was also considered and abandoned because it did not focus on the problem of real-time vs near-real-time sensing capacity was to replace all the 2.4GHz transceivers from the first generation parking nodes by 868MHz ones. However, we reached the conclusion that the additional costs of a provisional solution such as this one were not a price worth paying.

2) WIRELESS CHANNEL IMPAIRMENTS

Another unexpected transmission problem that we found in the city center was the existence of frequency jammers on the ISM bands around some private and public buildings. The reason behind this frequency jamming was security concerns. During the deployment location analysis we decided to focus on the city center and cover it with a dense multihop network. However, once deployed we realized that some devices were not reporting anything at all every now and then. After the corresponding analysis we discovered that the area in which these devices were installed was affected by a signal jammer that worked in an intermittent manner. This makes us to reconsider the deployment and move it out from that area.

C. EXTERNAL ENCLOSURE RELATED PRACTICALITIES

One of the most underestimated aspects of a real IoT deployment is the housing design phase. Bad decisions on this area can lead to IoT devices damage, therefore heavily increasing the total deployment cost.

The first idea a practitioner might typically adopt is to design a physically attractive enclosure. Although they are nice to showcase, custom designs are very expensive and

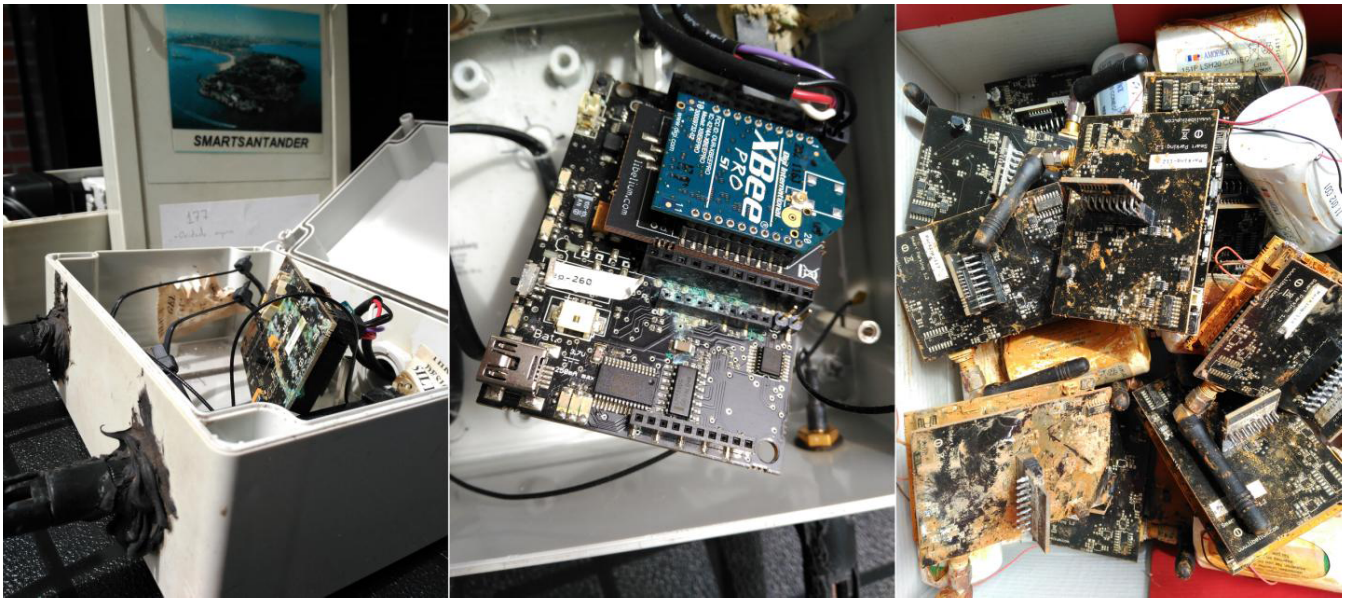

FIGURE 4. Examples of water-damaged IoT devices from SmartSantander deployment.

not as well tested on the field as standard ones. In addition, standard enclosures can often go unnoticed to the inexpert eye, so burglary will probably be diminished. However the most important factor that a practitioner need to evaluate is how well the enclosure behaves under outdoor environmental conditions.

In the Santander deployment case, we have experienced some unexpected problems that eventually led to the destruction of several nodes due to water leakage into the enclosure. Fig. 4 shows different examples of some of those devices.

The particularities of the different devices deployed as part of the SmartSantander testbed require different solutions. Still, we used IP67 certified standard enclosures in every use case. From a generic point of view, there are two fundamental use cases to consider:

1) HOUSING FOR NODES THAT CAN BE MANIPULATED AND REPAIRED ON THE FIELD

This is the case for most of the devices, which were deployed in lampposts and facades. This kind of nodes might just need a battery replacement or a quick adjustment and works on them do not have substantial impact on the usual city life.

The key requirement for this housing solution was clear: it has to allow in-place manipulation of the different modular elements inside it. Because of several reasons (e.g. radio propagation, minimize external interactions), these devices are mounted at least 3m above the ground, so the inside of the enclosure has to be easily accessible when using a ladder. Therefore, the enclosure needs to have a mechanism to open the cover, preferably without fully detaching it, while keeping the IP67 certification.

On the other hand, as most of these devices are connected to a public energy source we had to include a power protection switch (see subsection IV.A). On top of the already commented increase in price, it also had an important impact on the size of the external enclosure (an increase of around a 50%).

The last issue affecting the housing relates to the placement of transceivers antennas. In order to have higher antenna gains (5dBi vs 2dBi) that enhance the wireless transmission behavior of the nodes, we used antennas that would have required even larger housing if we were make them internal so they needed to be external. Existing recommendations at that time, indicated that using silicone rubber to keep water tightness, even if holes were made to the enclosure to place the needed antennas, was enough. This solution had been successfully tested by the device vendor, but again our experience after some time was not the same. The enclosure is water resistant against heavy rain in the short/mid period. Yet, water tend to form puddles on top of the enclosure and eventually drips into it, destroying the device.

The practical solution we adopted in order to be able to keep using already drilled enclosures without putting in risk the devices inside them, was to avoid having holes on the top of the enclosure when possible to prevent that water accumulation. In addition, another waterproof coating was also applied to re-seal them. Nevertheless, we decided to use internal patch antennas in all the new housing solutions and do not drill any hole on a certified enclosure. Our approach was to accept having a bigger enclosure able to accommodate all the components that to have to replace damaged devices. Fig. 5 shows the differences between the deployed patch antenna based housing solution and the previous one.

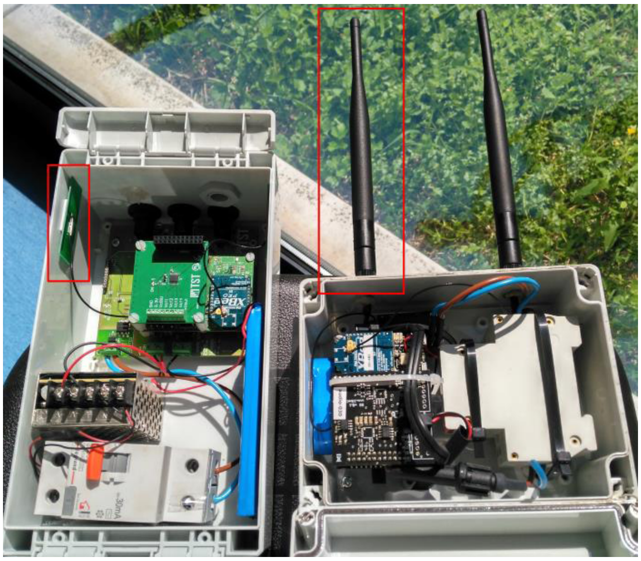

FIGURE 5. **Different antenna models: patch antenna (left) vs external antenna (right).**

2) HOUSING SOLUTIONS FOR NODES THAT NEED TO BE REPLACED AS A WHOLE ON THE FIELD, AND THEN REPAIRED ON THE LABORATORY

Working with this kind of nodes typically imply the need of civil work and/or traffic diversion, so interventions need to be as quick as possible. This is the case of all the devices buried under the asphalt, such as parking or traffic monitoring related ones.

Throughout the years we have tested several solutions which allowed battery replacement on the field while assuring devices integrity. We have learnt from practice that none of them satisfied this last part. Hence, on the long term, it is better to have hermetically sealed solutions (i.e. not accessible without destroying the enclosure) with acceptable battery duration.

The first iteration of parking nodes we deployed used a 10cm diameter cylindrical IP67 enclosure with a threaded cap. Even though it was buried in a hole in the asphalt, it was covered using a thin tar layer so that the replacement of their batteries could be done in around fifteen minutes by removing and re-applying the tar layer before placing the device back in the hole. However, this solution which in theory seemed sensible, showed several random problems with water getting into the enclosure due to condensation and humidity. We discover that due to temperature variations the seal gasket was not always working as expected, so the enclosure was not correctly sealed. Several different actions were taken to try to correct this problem: a) replacing the used tar by another material which was less breathable; b) using a different gasket; c) sealing the enclosure externally with a plastic material. As a result, although the problem was fixed, the time to replace a battery increased by a factor of 10, effectively turning the in-place battery replacement into a not viable solution.

When we had the opportunity of deploying a new generation of parking devices we chose a hermetically sealed solution. When a device fails we need to replace it as a whole, which is more expensive than just replacing a battery. However, the total costs when we take into consideration the needed time to actually replace a device on the field (with associated civil work and/or traffic diversion often involved) is comparable or even lower. Fig. 6 shows the differences between the aforementioned outdoor parking solutions.

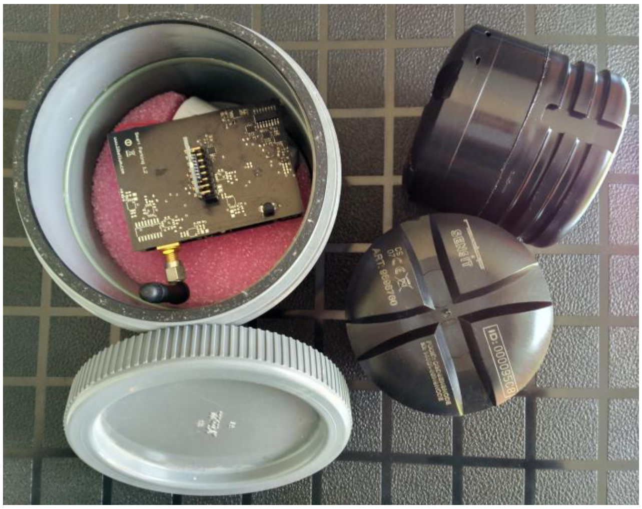

FIGURE 6. **Different generations of deployed outdoor parking solutions.**

As a conclusion for practitioners, carefully design outdoor housing solutions, they end up being almost as important as their content. Do not let a bad enclosure ruin your outdoor infrastructure. IoT devices can and will fail, so the easiest an on-the-field intervention can be the better. In fact, try to avoid those interventions as much as possible. When possible, incorporate a remote mechanism to cold-reset the nodes via hardware. Just in case this mechanism fails, consider the possibility of adding a transparent mechanism to cold-reboot the device without the need of opening the enclosure (e.g. using NFC or magnetic devices).

D. VENDOR DEPENDENCY RELATED PRACTICALITIES

Every smart city sensing infrastructure is typically deployed with a specific set of requirements in mind. Different vendors can fulfil these criteria using different approaches. However our experience shows that in most cases functionality prevails over interoperability. This is not bad per se, and it is often a good option on small scale deployments. Vertical solutions are usually well tested and the vendor can even provide a management platform, so it can speed up the time-to-market. Nevertheless, this is not always a good option when deploying massive IoT platforms. As it has been shown, they involve heterogeneous technology from different vendors, so the lack of a unified solution to manage the whole infrastructure can be highly inefficient. Having different silos for different use cases is not sustainable in the long term, as we might acquire dependency from the different involved vendors to be able to continue providing a city service.

The experience we have had with different vendors revealed that integrations of their systems with external platforms (such as SmartSantander [51] or FIWARE)[2] are not often well supported. In most cases we have had to implement specific polling solutions to periodically retrieve sensor information, which is far from being a scalable approach and also increases management complexity. Although this situation is understandable as infrastructure vendor strategy obey to market based rules, any infrastructure manager should avoid potential vendor lock-in situations whenever possible. In this sense, it might be interesting to consider those interoperability integrations as one extra requirement.

Another factor to balance is the openness of the solution. Even though it might offer a good cost-service ratio, obscuring most of the service layers and only being involved on the lowest and highest ones (deployment and result gathering process) presents disadvantages when something is not working as supposed. Longer delays to identify the problem, lack of proper technical information and extra costs are the typical scenario in this case. On the contrary, getting a deep knowledge of the whole stack is also not sensible approach as it might require a lot of resources that the infrastructure manager does not want or is not skilled to consume/provide. In SmartSantander we have had experiences at both ends of the spectrum, and the recommendation is to always look for open solutions where it is easier to find the trade-off between control over the infrastructure and the real obtained benefit.

Very specific requirements, and in particular those related to low level behavior, can jeopardize vendor copyrights. When this happens these requirements are not easy to fulfil. Off-the-shelf IoT devices are usually designed with a narrow functionality in mind so complexity is usually low. Even with a vendor agreement to modify low level functionalities, adding extra functionalities to a commercial solution has the potential risk of affecting the original one. In fact, depending on the agreement conditions, the vendor might not be held responsible of the possible problems, so the infrastructure owner would be the one in charge of fixing them. In an environment such as a smart city, where any physical intervention has a non-negligible cost, any modification need to be well tested in advance.

V. SMART CITY MANAGEMENT AND MONITORING IN PRACTICE

As it has been previously shown, SmartSantander network is particularly heterogeneous in terms of devices. In order to deal with such an infrastructure we have followed different approaches throughout the years.

At first, we decided to split the infrastructure into different domains in order to take advantage of the specificities of each kind of resource. This solution was very convenient to accelerate the developments on top of the infrastructure, hence to reduce the time-to-market. However, as a counterpart, it dramatically increased the infrastructure management complexity. In fact, as the number of resources and heterogeneity grew, the whole platform became what is usually referred as a ‘‘monster with thousand heads’’.

The aforementioned management scalability problems headed us into defining an information model which allowed us to focus on what is really important in an IoT scenario which is data and thus getting rid of the complexity of managing heterogeneous devices. However, homogeneous modelling of data is just half of the solution. It was also necessary to define a common functional modelling to inject (and also consume) the data generated into the SmartSantander platform. Once all the information in the platform was modelled following a common pattern, monitoring of the whole platform became easier and we were able to hide the underlying complexity from the SmartSantander core components. This has allowed us to turn the previous ad-hoc integration approximation into a much more generic one. As a result, external infrastructure providers are able to reuse the same set of existing modules and the platform managers do not need to be deeply involved on every integration project besides basic support.

This has an important implication in terms of reuse of pre-existing city owned infrastructure. Cities usually have different deployed infrastructures that are used to enhance different urban services. However, most of the time these systems are completely independent and managed in a ‘‘silo’’ approach. Integration of all this city infrastructure into a single, generic and accessible one provides a great example of how a platform such as SmartSantander contributes to ‘‘smartify’’ a city.

Information[3][4] and functional[5][6] models are nowadays being standardized. This section delves into the common information and functional modelling used in SmartSantander, which are in the root of some of these standardization works, briefly introducing the main concepts behind them. Afterwards, we focus on the monitoring system we have designed.

A. TERMINOLOGY AND DEFINITIONS

Before delving any deeper into the SmartSantander common information model specification, it is important to introduce some of the concepts and terms we used as a baseline for our design:

- A *SmartSantander resource* is any IoT device that is part of the infrastructure and, as such, produces observations or is able to receive actuation commands.

[2] FIWARE Smart Cities. [Online]. Available: https://www.fiware.org/smart-cities/, Accessed on: May, 30, 2017.

[3] ETSI Industry Specification Group on Context Information Management. [Online]. Available: https://portal.etsi.org/tb.aspx?tbid=854&SubTB=854, Accessed on: May, 30, 2017.

[4] FIWARE Data Models. [Online]. Available: https://www.fiware.org/data-models/, Accessed on: May, 30, 2017.

[5] NGSIv2 API documentation. [Online]. Available: http://docs.orioncontextbroker.apiary.io/# Accessed on: May, 30, 2017.

[6] OneM2M standard. [Online]. Available: http://www.onem2m.org/technical/published-documents, Accessed on: May, 30, 2017.

- *Service Experimentation (SEL)* consists of running experiments and/or applications based on the data gathered by SmartSantander sensor infrastructure and stored in a shared repository. Therefore, offering Sensing as a Service (SaaS) paradigm.
- *Capability* of an IoT device regards the sensing or actuating features of the devices. In this sense, we will talk about capabilities referring to the sensors or actuators with which a device is equipped.
- A *Topic* is the application domain to which a particular resource serves. One resource may serve to multiple topics.
- *External Infrastructure Provider* refers to any person or company which would like to expose their own resources through the SmartSantander facility.

B. COMMON INFORMATION MODELLING

The most important concept we apply to our information modelling is separation of concerns. In our experience, isolating resource model from observation model, while keeping a common taxonomy, helps to simplify management. IoT is rooted on devices but its real value is on the data that these devices can generate. Thus, it is necessary to manage devices but it is of utmost importance to allow extracting the value from the observations that these devices are continuously generating. In addition, keeping the static parameters (associated to the device) all together within the resource model and avoid the continuous repetition on all the observations helps to reduce the size of the whole dataset.

1) RESOURCE MODEL

Our resource information model is structured around five categories which gather the different features necessary to describe any IoT resource, with only a reduced set of mandatory parameters. These are identification, management, location, description and service; although only identification and management are mandatory categories.

- *Identification* category hosts the minimum identification details for that resource. In addition, it may contain a list of the different topics the resource belongs to.
- *Management* property includes information on resource's status and a list of management events the resource has gone through. Its content may be dynamically modified upon detection of changes on the resource status.
- *Location* property is modelled using GeoJSON schema. This parameter is not mandatory as mobile devices are not described by its location (the location information is however included as part of the observations model). Still, it is a highly recommended parameter in the case of fixed devices.
- *Description* category gathers a variety of descriptive information. In general, information included here might be mainly useful for humans although it can also be used as placeholder for semantic annotations.
- *Service* property contains the array of the capabilities that a resource provides. These capabilities will define the information that a device is able to produce, including the physical phenomena it can measure. Each capability is described by two compulsory attributes: phenomenon, which relates to the physical parameter the resource is able to sense or actuate; and uom, which relates to the unit of measurement used by the resource sensor. The possible values of these attributes are defined in a taxonomy that is part of the information model itself (cf. section V.B.3). Moreover, additional metadata (such as accuracy, frequency, latency, measurement range, precision, resolution, response time and sensitivity) associated with the capability can also be included.

2) OBSERVATION MODEL

An observation is defined as a set of sensor measurements that are generated by a single device at the same moment and on the same geographic location. As mentioned before, observations constitute an unbounded collection, so it is important to reduce the amount of information included to a minimum.

The dynamic parameters we consider as part of this minimum subset are: the originating device identifier, the observation timestamp, its location and a list of the different measurements that have been observed by the device. Each of these measurements is defined as one capability together with the specific measurement value. Although other parameters, such as additional metadata, could have been also included to fully describe the observation, we decided not to do it. The main reason behind it lies on the fact that all that information is statically linked to the capability on the corresponding resource description. This imposes the need of an extra indirection step to resolve the whole observation context, but on the other hand it lowers the infrastructure management burden while keeping a reasonable degree of context awareness on the observation model.

Although every observation inside the SmartSantander platform is comprised by all the aforementioned fields, an observation can be injected without some of those fields. In particular, timestamp and location properties can be omitted under certain circumstances. By allowing this, some low complexity IoT devices can still be compatible with the SmartSantander platform.

On the one hand, time synchronization is not mandatory. Although the recommendation is to always include a timestamp as close to the actual sampling time as possible, IoT devices can be too simple to have an internal clock, a GPS module, access to an NTP server or any other time-synchronization mechanism. In this sense, when timestamp is omitted the observation is automatically time-stamped by the system on injection. Of course, non-negligible drawbacks are the lack of a proper time zone (CET/CEST is used for convenience), an indeterminate lack of precision and the impossibility of measuring communication delays. Still, if these features are not part of mandatory application requirements, a lot of difficulties can be avoided by dropping time awareness features on a node.

On the other hand, location attribute can be omitted from observations when the device is a fixed one and that device's resource description includes a location property. In this case, geo-location of observations can be inferred from devices' descriptions.

3) TAXONOMY DEFINITION

A new domain specific taxonomy has been defined to denote physical parameters and units of measurement used in the SmartSantander scope. This dictionary has been created using already existing vocabularies or ontologies as a basis where possible. In fact, different external taxonomies are connected to the SmartSantander one to increase the interoperability (e.g. FIESTA-IoT[7] or FIWARE taxonomies). Besides, in order to make them easier to adopt by application developers, every entry includes a human-readable textual definition and the whole lexicon, together with the existing relations between entries, can be accessed freely from the SmartSantander APIs. Moreover, the taxonomy is considered extendable as it can allocate new phenomenon (or unit of measurement) not yet observed. Once a request for addition to the taxonomy is validated by an administrator, it will be promoted to the official list and the platform will be able to accept new capabilities.

C. COMMON FUNCTIONAL MODELLING

Information modelling is just one of the two main aspects in data management. As part of SmartSantander SEL value proposition, the functional modelling (i.e. the way to access data) has been redefined to be more user oriented. In this sense, instead of abstracting experimenter's interactions with our platform, we made them more focused on actual data consumer needs. In our opinion, just offering sensor information based on resource identifiers does not provide an appropriate user experience. However, being able to answer natural queries such as "I'd like all the observations containing temperature measurements above 10 degrees Celsius in the area defined by this polygon" are a much more user-friendly option, even if the queries still need to be formatted in a domain specific JSON based format.

Besides, separation of concerns has also been applied into the definition of the common functional modelling. On the one hand, functional roles that the different stakeholders around an IoT scenario can play have been clearly delimited:

- An *infrastructure provider* uses the platform to publish its sensor data.
- A *data consumer* gets sensor information from the platform and provides added value services on top of it.
- A *data prosumer* (i.e. combination of both roles) is also considered. Raw sensor data can be aggregated (or any other way of data transformation) to generate virtual sensors based on it. Then he can register those virtual sensors in the platform under a new domain and inject the generated virtual observations for others to use.
- *Platform administrators* manage the whole platform core infrastructure and can provide support when needed.

On the other hand, as different roles' demands towards the smart city platform are not homogeneous, four separate subsystems are also defined. They cover different functionalities within a smart city platform:

- The *security subsystem* is responsible for the authentication, authorization and management of the users in the system.
- The *resource subsystem* handles resource description information and takes part on the resource registration process.
- The mission of the t*axonomy subsystem* is the definition of the accepted vocabulary associated with the different sensor parameters supported
- The *information subsystem* is in charge of adapting, storing and exposing sensor observations generated by the different resources. Data from this subsystem can be retrieved both in synchronous and asynchronous mode.

Access to all these subsystems is provided through a cohesive RESTful API. Albeit it follows a request-response communication scheme it can also be used in the context of asynchronous publication/subscription services. In this latter case, experimenters can use such an interface to synchronously manage the configuration of asynchronous notifications.

D. INFRASTRUCTURE MONITORING

Monitoring a heterogeneous IoT infrastructure such as the one SmartSantander offer is heavily influenced by how well organized the platform management layer is. Well-defined information and functional models, together with modular management platform architecture can lead to easier solutions. Moreover, the usage of independent management platforms or different data paths, often observed when adopting "silo" approaches, usually require the introduction of extra translation layers that can cause difficulties on the monitoring process.

Multiple infrastructure monitoring aspects can be pursued as part of the IoT infrastructure management efforts but in our experience, key ones are device status and data quality monitoring. The first one is directly linked with the resource subsystem and is considered a crucial element for infrastructure maintenance. The second one is typically associated with the information subsystem and it is of great interest for providing meaningful meta-data to data consumers. In addition, device status can also be inferred based on information extracted from data quality analysis so data quality monitoring is also used for infrastructure maintenance purposes.

Next subsections summarizes the solutions adopted for these two areas, which are also represented in Fig. 7.

[7] FIESTA-IoT Ontology. [Online]. Available: http://ontology.fiesta-iot.eu/ontologyDocs/ Accessed on: May, 30, 2017.

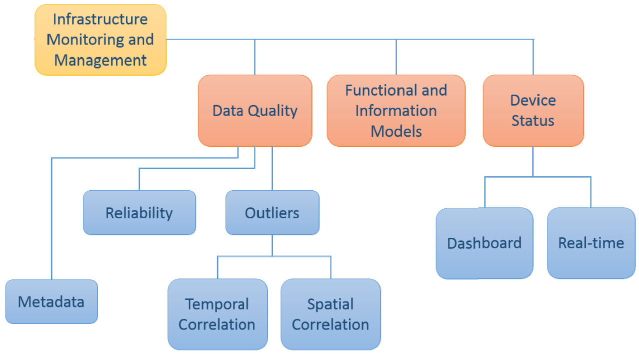

FIGURE 7. **Addressed management & monitoring related challenges.**

1) DATA QUALITY MONITORING

Among the different techniques that can be applied on this specific task, we have focused on outlier analysis and data cleansing. Both techniques are applied as the information is injected into the core management platform. At this point all observations are time-stamped and geo-located, so there is enough context information to execute data analytics mechanisms over it. By doing this, we avoid the need of in-network intermediate computations as network status self-awareness is not required. Both techniques are deeply related and are concurrently applied on a per-phenomena basis.

Before starting with more complex calculations all out-of-range (where range is dynamically obtained from temporal and spatial correlation) measurements are tagged. We have observed that problems with sensors are mostly due to two factors: either the sensor connection to the host node fails or the sensor itself get stuck on a single value. By applying this straightforward data cleansing technique we almost eliminate one entropy source on the outlier analysis methodology. All these measurements are not discarded from the system, but just tagged (i.e. through adding meta-data attributes to the observation and/or resource description) in order to enable further analysis.

After that, outlier analysis is applied by including the neighbor measurements in a vicinity area defined both by the distance (i.e. spatial correlation) and the number of observations (i.e. temporal correlation) in that cluster. We conveniently weight each value with the distance, so nearest (in Euclidean distance and time) ones will have bigger impact on the decision. In addition, different weights are also assigned based on the quality of the sensor. Some sensors have better accuracy figures so they are better references to analyze possible outlier values. As all the information is structured following the previously described observation model, all measurements are treated equally, whether they are generated by fixed or mobile devices.

It is important to mention that our data analysis was initially only focused on spatial correlation (i.e. analyzing neighboring nodes datasets to look for outliers). However, we soon realized that spatial correlation alone do not always offer good enough overall results for all the phenomena set. In general, environmental related phenomena present sufficient spatial stability, hence are good candidates for applying this kind of techniques exclusively. On the contrary, there are several use cases where no direct correlation exists between adjacent nodes (e.g. presence detection related phenomena) and spatial analysis is not enough to infer wrong values. In fact, some of the spatial correlated phenomena can also experiment large fluctuations within small distances due to external factors. This is the case of outdoor/indoor neighbors, or luminosity nodes placed on shadow areas during daylight or next to a lamppost during night. The combination of spatial correlation techniques together with temporal correlation ones can help to identify whether or not a node is providing good quality data. In this sense, temporal correlation of the measurements provided by the node and their neighbors is also applied.

Finally, once outliers are identified data cleansing mechanism is re-applied. This mechanism filters out all tagged measurements and injects them into a different stream. This data is then processed by different platform modules which produce inputs to the device status monitoring system. In addition, data is also available to external consumers, as it is often interesting for extracting different failure patterns or metrics (e.g. time between failures).

2) DEVICE STATUS MONITORING

Device status monitoring do not work on top of information layer, but at resource level. It is in charge of monitoring the physical infrastructure and generate alarms when a node do not behaves as expected.

Using inputs from the data quality monitoring, together with real-time analysis of the received streams, it infers the actual status of every deployed device. When an unexpected event is detected (e.g. a node does not report observations during the expected time window or it produces outliers) the needed modifications are performed on the corresponding resource description (specifically on the management category). Reports are then generated every day and sent to the IoT deployment managers, so they can coordinate actuations on the infrastructure. Moreover, the information on the management category of every resource description is graphically presented on the infrastructure manager dashboard for immediate grasp of infrastructure overall status.

VI. CONCLUSIONS

IoT deployments in general and in smart cities in particular are complex scenarios. While solutions in this area tend to be focused on the "things", management of the data produced by the infrastructure, where the actual value of these deployments is, requires the same, or even more attention.

This paper has presented the solutions adopted for a real-world IoT-enabled smart city deployment in the city of Santander. These solutions are grouped around both, the hardware aspects and actual in-site deployment of the IoT devices, and the monitoring of these devices and the data that they are continuously producing.

The main contribution from this paper is presenting practical solutions together with the motivations and lessons learnt while they were developed and put in place. In this sense, we have described the hidden pitfalls that we encountered during our own experience throughout the deployment and management processes. To the best of our knowledge, there are very little information available on IoT infrastructure for smart cities deployment and management best-practices. Thus, we believe that the description provided in this paper can be useful for practitioners in their own deployments.

As part of the future work we are exploring the use of more sophisticated data analytics (mainly Principal Component and Artificial Neural Networks) which could provide even better insights on the data quality and what is more important support in-advance management of deployed devices.

REFERENCES

[1] L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A survey," *Comput. Netw.*, vol. 54, no. 15, pp. 2787–2805, Oct. 2010.

[2] A. Rachedi, M. H. Rehmani, S. Cherkaoui, and J. J. P. C. Rodrigues, "IEEE access special section editorial: The plethora of research in Internet of Things (IoT)," *IEEE Access*, vol. 4, pp. 9575–9579, 2016.

[3] Z. Sheng, S. Yang, Y. Yu, A. Vasilakos, J. McCann, and K. Leung, "A survey on the IETF protocol suite for the Internet of Things: Standards, challenges, and opportunities," *IEEE Wireless Commun.*, vol. 20, no. 6, pp. 91–98, Dec. 2013.

[4] A. Al-Fuqaha, M. Guizani, M. Mohammadi, M. Aledhari, and M. Ayyash, "Internet of Things: A survey on enabling technologies, protocols, and applications," *IEEE Commun. Surveys Tuts.*, vol. 17, no. 4, pp. 2347–2376, 4th Quart., 2015.

[5] P. Neirotti, A. De Marco, A. C. Cagliano, G. Mangano, and F. Scorrano, "Current trends in smart city initiatives: Some stylised facts," *Cities*, vol. 38, pp. 25–36, Jun. 2014.

[6] L. Zhang *et al.*, "A remote medical monitoring system for heart failure prognosis," *Mobile Inf. Syst.*, vol. 2015, Sep. 2015, Art. no. 406327.

[7] S. M. Riazul Islam, D. Kwak, M. Humaun Kabir, M. Hossain, and K.-S. Kwak, "The Internet of Things for health care: A comprehensive survey," *IEEE Access*, vol. 3, pp. 678–708, Jun. 2015.

[8] K. Zheng, S. Zhao, Z. Yang, X. Xiong, and W. Xiang, "Design and implementation of LPWA-based air quality monitoring system," *IEEE Access*, vol. 4, pp. 3238–3245, 2016.

[9] K.-L. Tsai, F.-Y. Leu, and I. You, "Residence energy control system based on wireless smart socket and IoT," *IEEE Access*, vol. 4, pp. 2885–2894, 2016.

[10] J. Wan, M. Yi, D. Li, C. Zhang, S. Wang, and K. Zhou, "Mobile services for customization manufacturing systems: An example of industry 4.0," *IEEE Access*, vol. 4, pp. 8977–8986, 2016.

[11] A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, "Internet of Things for smart cities," *IEEE Internet Things J.*, vol. 1, no. 1, pp. 22–32, Feb. 2014.

[12] J. M. Hernández-Muñoz *et al.*, "Smart cities at the forefront of the future Internet," in *The Future Internet Assembly*. Berlin, Germany: Springer, 2011, pp. 447–462.

[13] H. Chourabi *et al.*, "Understanding smart cities: An integrative framework," in *Proc. 45th Hawaii Int. Conf. Syst. Sci.*, Jan. 2012, pp. 2289–2297.

[14] Z. Liu, M. Wu, K. Zhu, and L. Zhang, "SenSafe: A smartphone-based traffic safety framework by sensing vehicle and pedestrian behaviors," *Mobile Inf. Syst.*, vol. 2016, Sep. 2016, Art. no. 7967249.

[15] R. Petrolo, V. Loscrì, and N. Mitton, "Towards a smart city based on cloud of things, a survey on the smart city vision and paradigms," *Trans. Emerg. Telecommun. Technol.*, vol. 28, no. 1, p. e2931, Jan. 2017.

[16] W. M. da Silva *et al.*, "Smart cities software architectures: A survey," in *Proc. 28th Annu. ACM Symp. Appl. Comput. (SAC)*, 2013, pp. 1722–1727.

[17] A. Gluhak, S. Krco, M. Nati, D. Pfisterer, N. Mitton, and T. Razafindralambo, "A survey on facilities for experimental Internet of Things research," *IEEE Commun. Mag.*, vol. 49, no. 11, pp. 58–67, Nov. 2011.

[18] J. H. Lee, M. G. Hancock, and M.-C. Hu, "Towards an effective framework for building smart cities: Lessons from Seoul and San Francisco," *Technol. Forecast. Soc. Change*, vol. 89, pp. 80–99, Nov. 2014.

[19] S. Zygiaris, "Smart city reference model: Assisting planners to conceptualize the building of smart city innovation ecosystems," *J. Knowl. Econ.*, vol. 4, no. 2, pp. 217–231, Jun. 2013.

[20] T. Bakici, E. Almirall, and J. Wareham, "A smart city initiative: The case of barcelona," *J. Knowl. Econ.*, vol. 4, no. 2, pp. 135–148, Jun. 2013.

[21] T. Gea, J. Paradells, M. Lamarca, and D. Roldán, "Smart cities as an application of Internet of Things: Experiences and lessons learnt in barcelona," in *Proc. 7th Int. Conf. Innov. Mobile Internet Services Ubiquitous Comput.*, Jul. 2013, pp. 552–557.

[22] L. Sanchez *et al.*, "SmartSantander: IoT experimentation over a smart city testbed," *Comput. Netw.*, vol. 61, pp. 217–238, Mar. 2014.

[23] C. Adjih *et al.*, "FIT IoT-LAB: A large scale open experimental IoT testbed," in *Proc. IEEE 2nd World Forum Internet Things (WF-IoT)*, Dec. 2015, pp. 459–464.

[24] M. Vögler, J. Schleicher, C. Inzinger, S. Nastic, S. Sehic, and S. Dustdar, "LEONORE—Large-scale provisioning of resource-constrained IoT deployments," in *Proc. IEEE Symp. Service-Oriented Syst. Eng.*, Mar./Apr. 2015, pp. 78–87.

[25] O. Vermesan and P. Friess, *Internet of Things Applications: From Research and Innovation to Market Deployment*. Aalborg, Denmark: River, 2014.

[26] P. Barnaghi, W. Wang, C. Henson, and K. Taylor, "Semantics for the Internet of Things: Early progress and back to the future," *Int. J. Semantic Web Inf. Syst.*, vol. 8, no. 1, pp. 1–21, 2012.

[27] J. Kiljander *et al.*, "Semantic interoperability architecture for pervasive computing and Internet of Things," *IEEE Access*, vol. 2, pp. 856–873, 2014.

[28] M. Blackstock and R. Lea, "IoT interoperability: A hub-based approach," in *Proc. Int. Conf. Internet Things (IOT)*, Oct. 2014, pp. 79–84.

[29] P. Desai, A. Sheth, and P. Anantharam, "Semantic gateway as a service architecture for IoT interoperability," in *Proc. IEEE Int. Conf. Mobile Services*, Jun./Jul. 2015, pp. 313–319.

[30] J. Lanza, L. Sanchez, D. Gomez, T. Elsaleh, R. Steinke, and F. Cirillo, "A proof-of-concept for semantically interoperable federation of IoT experimentation facilities," *Sensors*, vol. 16, no. 7, p. 1006, Jun. 2016.

[31] R. Lea, M. Blackstock, N. Giang, and D. Vogt, "Smart cities: Engaging users and developers to foster innovation ecosystems," in *Proc. ACM Int. Joint Conf. Pervasive Ubiquitous Comput., ACM Int. Symp. Wearable Comput. (UbiComp)*, 2015, pp. 1535–1542.

[32] L. Sánchez, I. Elicegui, J. Cuesta, L. Muñoz, and J. Lanza, "Integration of utilities infrastructures in a future Internet enabled smart city framework," *Sensors*, vol. 13, no. 11, pp. 14438–14465, 2013.

[33] M. Swan, "Sensor mania! The Internet of Things, wearable computing, objective metrics, and the quantified self 2.0," *J. Sens. Actuator Netw.*, vol. 1, no. 3, pp. 217–253, Nov. 2012.

[34] L. Gao and X. Bai, "A unified perspective on the factors influencing consumer acceptance of Internet of Things technology," *Asia–Pacific J. Marketing Logistics*, vol. 26, no. 2, pp. 211–231, Apr. 2014.

[35] A. Karkouch, H. Mousannif, H. Al Moatassime, and T. Noel, "Data quality in Internet of Things: A state-of-the-art survey," *J. Netw. Comput. Appl.*, vol. 73, pp. 57–81, Sep. 2016.

[36] Y. Zhang, N. Meratnia, and P. Havinga, "Outlier detection techniques for wireless sensor networks: A survey," *IEEE Commun. Surveys Tuts.*, vol. 12, no. 2, pp. 159–170, 2nd Quart., 2010.

[37] T. Dasu and T. Johnson, *Exploratory Data Mining and Data Cleaning*. Hoboken, NJ, USA: Wiley, 2003.

[38] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection," *ACM Comput. Surv.*, vol. 41, no. 3, p. 15, Jul. 2009.

[39] A. Patcha and J.-M. Park, "An overview of anomaly detection techniques: Existing solutions and latest technological trends," *Comput. Netw.*, vol. 51, no. 12, pp. 3448–3470, 2007.

[40] S. Sathe, T. Papaioannou, H. Y. Jeung, and K. Aberer, "A survey of model-based sensor data acquisition and management," in *Managing and Mining Sensor Data*. Boston, MA, USA: Springer, 2013, pp. 9–50.

[41] C. C. Aggarwal, N. Ashish, and A. Sheth, "The Internet of Things: A survey from the data-centric perspective," in *Managing and Mining Sensor Data*. Boston, MA, USA: Springer, 2013, pp. 383–428.

[42] J. Lei, H. Bi, Y. Xia, J. Huang, and H. Bae, "An in-network data cleaning approach for wireless sensor networks," *Intell. Autom. Soft Comput.*, vol. 22, no. 4, pp. 599–604, Oct. 2016.

[43] N. K. Thanigaivelan, E. Nigussie, R. K. Kanth, S. Virtanen, and J. Isoaho, "Distributed internal anomaly detection system for Internet-of-Things," in *Proc. 13th IEEE Annu. Consum. Commun. Netw. Conf. (CCNC)*, Jan. 2016, pp. 319–320.

[44] S. Gill and B. Lee, "A framework for distributed cleaning of data streams," *Procedia Comput. Sci.*, vol. 52, pp. 1186–1191, Jan. 2015.

[45] Y. Zhang, C. Szabo, and Q. Z. Sheng, *Cleaning Environmental Sensing Data Streams Based on Individual Sensor Reliability*. Cham, Switzerland: Springer, 2014, pp. 405–414.

[46] L. Sanchez *et al.*, "SmartSantander: The meeting point between future Internet research and experimentation and the smart cities," in *Proc. Future Netw. Mobile Summit*, Jun. 2011, pp. 978–985.

[47] R. van Kranenburg *et al.*, "Co-creation as the key to a public, thriving, inclusive and meaningful EU IoT," in *Ubiquitous Computing and Ambient Intelligence. Personalisation and User Adapted Services*. Cham, Switzerland: Springer, 2014, pp. 396–403.

[48] J. Lanza *et al.*, "Large-scale mobile sensing enabled Internet-of-Things testbed for smart city services," *Int. J. Distrib. Sensor Netw.*, vol. 11, no. 8, 2015.

[49] J. Lanza *et al.*, "Smart city services over a future Internet platform based on Internet of Things and cloud: The smart parking case," *Energies*, vol. 9, no. 9, p. 719, Sep. 2016.

[50] M. Foulonneau, S. Martin, and S. Turki, *How Open Data are Turned Into Services?*. Cham, Switzerland: Springer, 2014, pp. 31–39.

[51] J. Lanza *et al.*, "Managing large amounts of data generated by a smart city Internet of Things deployment," *Int. J. Semantic Web Inf. Syst.*, vol. 12, no. 4, pp. 22–42, 2016.

PABLO SOTRES received the Telecommunications Engineering degree from the University of Cantabria, Spain, in 2008. He has been involved in several different international projects framed under the smart city paradigm, such as SmartSantander; and related to inter-testbed federation, such as Fed4FIRE, Fed4FIRE+, and Wise-IoT. He is currently a Research Fellow with the Network Planning and Mobile Communications Laboratory, Communications Engineering Department, University of Cantabria.

JUAN RAMÓN SANTANA received the Telecommunications Engineering degree from the University of Cantabria in 2010. He was an Intern with the University of Strathclyde, Glasgow, involved in IoT solutions. He has been involved in several projects, such as SmartSantander, EAR-IT or FESTIVAL, and European collaborative projects related to the smart city paradigm and the Internet of Things. He is currently a Research Fellow with the Network Planning and Mobile Communications Laboratory, Telecommunication Research Group, University of Cantabria. Among his research interests are wireless sensor networks, M2M communications, and mobile phone application research.

LUIS SÁNCHEZ received the Telecommunications Engineering and Ph.D. degrees from the University of Cantabria, Spain, in 2002 and 2009, respectively. He is currently an Associate Professor with the Department of Communications Engineering, University of Cantabria. He is active on IoT-enabled smart cities, meshed networking on heterogeneous wireless scenarios and optimization of network performance through cognitive networking techniques. He has a long research record working on projects belonging to the 5th, 6th, 7th, and H2020 EU Framework Programs. He has authored more than 60 papers at international journals and conferences and co-authored several books.

JORGE LANZA received the Ph.D. degree in telecommunications engineering from the University of Cantabria, Spain, in 2014. He has participated in several research projects, national and international, with both private and public funding. He is currently a Senior Researcher with the Network Planning and Mobile Communications Laboratory, University of Cantabria. His current research focuses on IoT infrastructures toward federating deployments in different locations using semantics technologies. In addition, his work has included combined mobility and security for the wireless Internet.

LUIS MUÑOZ received the Telecommunications Engineering and Ph.D. degrees from the Polytechnical University of Cataluña, Spain, in 1990 and 1995, respectively. He is currently the Head of the Network Planning and Mobile Communications Laboratory, Communications Engineering Department, University of Cantabria, Spain. He has participated in several National and European research projects belonging to the 4th, 5th, 6th, 7th, and H2020 Framework Programs in which he was the Technical Manager of SmartSantander. He has authored over 150 journal and conference papers. His research focuses on advanced data transmission techniques, heterogeneous wireless multihop networks and applied mathematical methods for telecommunications. He serves as an editor of several journals. In parallel to this activity, he serves as a Consultant for the Spanish Government as well as for different companies in Europe and USA.