

Chapter 2

Building a Smart City Ecosystem for Third Party Innovation in the City of Heraklion



Manos Kalaitzakis, Manousos Bouloukakis, Pavlos Charalampidis, Manos Dimitrakis, Giannis Drossis, Alexandros Fragkiadakis, Irimi Fundulaki, Katerina Karagiannaki, Antonis Makrogiannakis, Georgios Margetis, Athanasia Panousopoulou, Stefanos Papadakis, Vassilis Papakonstantinou, Nikolaos Partarakis, Stylianos Roubakis, Elias Tragos, Elisjana Ymeralli, Panagiotis Tsakalides, Dimitris Plexousakis and Constantine Stephanidis

Abstract This paper describes the implementation of an Internet of Things (IoT) and Open Data infrastructure by the Institute of Computer Science of the Foundation for Research and Technology—Hellas (FORTH-ICS) for the city of

M. Kalaitzakis (✉) · M. Bouloukakis · P. Charalampidis · M. Dimitrakis · G. Drossis
A. Fragkiadakis · I. Fundulaki · A. Makrogiannakis · G. Margetis · A. Panousopoulou
S. Papadakis · V. Papakonstantinou · N. Partarakis · S. Roubakis · E. Tragos · E. Ymeralli
P. Tsakalides · D. Plexousakis · C. Stephanidis
Foundation for Research and Technology – Hellas (FORTH), Institute of Computer Science,
N. Plastira 100, Vassilika Vouton, 70013 Heraklion, Crete, Greece
e-mail: mkalaitz@ics.forth.gr

M. Bouloukakis
e-mail: mboulou@ics.forth.gr

P. Charalampidis
e-mail: pcharala@ics.forth.gr

M. Dimitrakis
e-mail: mdimitrak@ics.forth.gr

G. Drossis
e-mail: drossis@ics.forth.gr

A. Fragkiadakis
e-mail: alfrag@ics.forth.gr

I. Fundulaki
e-mail: fundul@ics.forth.gr

A. Makrogiannakis
e-mail: makrog@ics.forth.gr

G. Margetis
e-mail: gmarget@ics.forth.gr

© Springer Nature Switzerland AG 2019

A. Stratigea and D. Kavroudakis (eds.), *Mediterranean Cities and Island Communities*, Progress in IS, https://doi.org/10.1007/978-3-319-99444-4_2

19

Heraklion, focusing on the application of mature research and development outcomes in a Smart City context. These outcomes mainly fall under the domains of Telecommunication and Networks, Information Systems, Signal Processing and Human Computer Interaction. The infrastructure is currently being released and becoming available to the municipality and the public through the Heraklion Smart City web portal. It is expected that in the future such infrastructure will act as one of the pillars for sustainable growth and prosperity in the city, supporting enhanced overview of the municipality over the city that will foster better planning, enhanced social services and improved decision-making, ultimately leading to improved quality of life for all citizens and visitors.

Keywords IoT · Smart cities · Open data · Data analytics · Smart city visualization · Sustainable growth · Third party innovation

2.1 Introduction

As the IoT landscape is growing, expectations are also raised from both the citizens (i.e. consumer) and the industry's (i.e. producer) point of view. The smart city concept aims at combining modern technology with social activities in the city to address societal challenges. The term itself has been defined in many different

A. Panousopoulou
e-mail: apanouso@ics.forth.gr

S. Papadakis
e-mail: stefpap@ics.forth.gr

V. Papakonstantinou
e-mail: papv@ics.forth.gr

N. Partarakis
e-mail: partarak@ics.forth.gr

S. Roubakis
e-mail: roub@ics.forth.gr

E. Tragos
e-mail: etragos@ics.forth.gr

E. Ymeralli
e-mail: ymeralli@ics.forth.gr

P. Tsakalides
e-mail: tsakalid@ics.forth.gr; tsakalid@csd.uoc.gr

D. Plexousakis
e-mail: dp@ics.forth.gr; dp@csd.uoc.gr

C. Stephanidis
e-mail: cs@ics.forth.gr; cs@csd.uoc.gr

K. Karagiannaki · P. Tsakalides · D. Plexousakis · C. Stephanidis
Department of Computer Science, University of Crete, Heraklion, Crete, Greece
e-mail: karayan@csd.uoc.gr

ways; and several working definitions have been put forward and adopted for both practical and academic use (Chourabi et al. 2012). According to Caragliu et al. (2011), the common characteristics of smart cities are the: (a) utilization of a networked infrastructure for improving economic and governance efficiency and enabling social, cultural and urban development; (b) underlying emphasis on business-led urban development; (c) usage of high-tech and creative industries in long-run urban growth, and (d) pursuit of social and environmental sustainability. For the purposes of the research work reported in this paper, the following broad definition is adopted:

A city striving to make itself “smarter” (more efficient, sustainable, equitable, and livable)
(Natural Resources Defense Council n.d.)

Worldwide, an enormous amount of resources is currently dedicated to the development of smart city ecosystems. Additionally, in such contexts, the widespread usage of Open Data addresses the need of data availability and access to knowledge. These two notions can be combined in a single vertical platform that serves a twofold purpose, namely the: monitoring, harvesting and analyzing of data flows on the one hand, and dissemination of these data to all interested parties on the other.

From a technical perspective, the smart city remains a challenge due to the lack of interoperability of the heterogeneous technologies currently used in city and urban development. In this respect, the IoT vision, in a smart city context, aims at providing a single extendable and interoperable smart platform that can ultimately become a building block to realize a unified urban-scale ICT platform (Chourabi et al. 2012). As such it can unleash the potential of the smart city vision (Hernández-Muñoz et al. 2011; Mulligan and Olsson 2013).

Taking into account the aforementioned consideration, this paper presents the challenges faced in designing and implementing an IoT and open data infrastructure for the Municipality of Heraklion, Crete. The infrastructure is being developed under a Programmatic Agreement between the Municipality of Heraklion and FORTH-ICS, one of the largest research centers in Greece. In this process, R&D outcomes of four laboratories of FORTH-ICS are exploited, including research outcomes achieved through the collaboration of FORTH-ICS and the Municipality of Heraklion in the context of European funded projects [e.g. FP7 RERUM project (RERUM n.d.)]. The discussion on this paper focuses both on fundamental technological components and strategic decisions towards an open innovation ecosystem.

Currently, the IoT infrastructure consists of the following nodes: (a) environmental and weather monitoring, (b) air quality monitoring, (c) water quality and management monitoring, and (d) smart parking. Furthermore, the intelligence of the smart city platform is provided through an open data middleware and portal, a data analytics platform and an open API for third party innovation. By releasing this open infrastructure, it is expected that the city of Heraklion will, in the near future, be capable of: presenting a new business model for sustainable growth, based on cutting edge IoT technology; and promoting more intelligent city data collection,

analysis, and rationalization, so as to foster intelligent decision-making by the city's authorities. This will also allow SME's to exploit and create added-value services for both citizens and visitors, including the elderly and people with disabilities.

2.2 Background—State-of-the-Art

This section focuses on the technological advancements that are related to the main architectural components of a smart city infrastructure. These range from an overview of the state of the art in IoT and smart cities to the different approaches regarding network architectures, IoT and Open Data middleware technologies, data analytics and information visualisation.

2.2.1 *State-of-the-Art in IoT and Smart Cities*

The evolution of computing technology is inevitably leading to a new ICT landscape, where in a near future the objects of everyday life will be equipped with microcontrollers, transceivers for digital communication, and suitable protocol stacks that will make them able to communicate with one another and with the users, becoming an integral part of the Internet (Atzori et al. 2010). Such a technological evolution is already happening and expands to the city context, as it responds to the strong interest of many national governments and regional organizations (e.g. municipalities) to adopt ICT solutions in the management of public affairs, thus realizing the so-called smart city concept (Chourabi et al. 2012).

In smart cities, IoT technology has the main role in supporting value-added services for the administration of the city and its citizens (Zanella and Vangelista 2014). At the same time, the existence of a connected, distributed and autonomous sensing infrastructure leads to an exponential growth of data that should be collected, maintained, processed and reasoned upon (Molinari et al. 2014).

Taking into account this situation, building a Smart City ecosystem requires special care. The main challenge is to ensure that the infrastructure will be scalable, resilient to changes, secure from the point of view of both the infrastructure and users, and open for third party innovation in the city, by exposing functionality and data.

In the context of this case study, in order to address the aforementioned challenges, five important components of the infrastructure were carefully designed, so as to ensure that the ecosystem will be able to support the today's requirements and at the same time to evolve as new technologies and requirements are generated. These components are described in detail in Sect. 2.4 and they consist of the:

(i) physical sensor infrastructure and the network communication protocols, (ii) IoT middleware, (iii) Open Data middleware, (iv) analytics component and (v) visualization layer.

2.2.2 Network Topologies/Architectures for Smart City Infrastructure

IoT refers to the interconnection of a large number of smart objects that are mainly able to sense the environment and report their findings to centralized entities like cloud servers. The term smart object refers to entities like wireless sensors, smart phones, smart cars, etc. Due to the rapid advance of technology in terms of hardware, miniature wireless sensors can provide a plethora of sensory data like ambient temperature, humidity, gas concentration, weather monitoring, etc. Moreover, the proliferation of operating systems for miniature devices like the ContikiOS (Contiki n.d.), TinyOS (n.d.), etc., has enabled the interconnection of these devices using IP-based protocols (e.g. IPv6). On top of IP, several other communication protocols provide energy-efficient interconnectivity between the sensors and the backend servers.

For efficient data collection, a number of protocols are used like the COAP, MQTT, etc. COAP (Constrained Application Protocol) is a network-oriented protocol based on REST architecture and it executes over UDP to avoid congestion in Wireless Sensor Networks (WSNs). Resources in COAP are identified by URIs, and clients send requests to a server based on a specific URI. Clients interested in a specific resource (sensory data), subscribe to a server (e.g. Gateway) and receive the corresponding measurements as soon as these become available.

From a telecommunication and networks point of view, IoT architectures currently face a number of challenges:

- *Security*: Protocol inefficiencies and software vulnerabilities have led to numerous attacks against IoT networks, like DoS attacks, wormholes, Sybil attacks, routing attacks, eavesdropping, fabrication, message replay, etc.
- *Privacy*: As sensors can collect and convey sensitive information, as for example in the case of healthcare scenarios, privacy preservation is of paramount importance. IoT wider acceptance also depends on privacy preservation.
- *Interoperability*: A large number of IoT architectures have been proposed so far, with significant contributions and technological advances in several areas (e.g. energy efficiency, etc.). However, interoperability has not been properly addressed, leading to ‘silos’ where each single architecture is isolated from the others.
- *Scalability*: Very often, the smart objects (i.e. sensors) are severely constrained devices in terms of processing, memory and storage. This, along with several protocol inefficiencies, degrades IoT networks’ performance in terms of throughput, delay and packet loss.

A few “lighthouse” projects have significantly contributed to the maturity of IoT. These are:

- *RERUM* (RERUM n.d.): a successful IoT platform, funded by EC as an FP7 project. The functional architecture of RERUM is based on the architectural reference model of IoT-A. However, it follows not only a service-oriented approach, like IoT-A as well as most IERC projects (Tragos et al. 2016), but also assumes that the devices have an important role in ensuring the security and privacy of the architecture.
- *IoT-A*: one of the first IoT projects, focused on the creation of a generic Architectural Reference Model (ARM) that is used for deriving concrete IoT architectures. The ARM consists of several sub-models that set the scope of the IoT design space (Carrez 2013). These are the: (i) IoT communication model, (ii) IoT trust, security and privacy model, (iii) IoT functional model, (iv) IoT information model, and (v) IoT domain model.
- *FIWARE*: aims at building a core platform for the Future Internet and adopts the notion of “Generic Enablers”, components that offer reusable and commonly shared functions, serving a number of application scenarios. The generic enablers provide architecture reference model for: (i) cloud hosting, (ii) data/context management, (iii) applications/services ecosystem and delivery framework, (iv) IoT services enablement, (v) interface to network and devices, and (vi) security (Krcro et al. 2014).

IoT technology fragmentation, along with the lack of global IoT standards, has led to isolated IoT systems, incapable of communicating with other systems that use different technologies, thus creating barriers for interoperable IoT systems. The *INTER-IoT* project (INTER-IoT Project n.d.) addresses interoperability issues by providing all those building blocks needed in order this to be achieved, including a framework, a methodology and the associated APIs and tool boxes.

2.2.3 *IoT Middlewares*

Typical IoT architectures are based on a number of sensors, often heterogeneous in nature, with different software and hardware capabilities. Moreover, several operations, like service discovery and orchestration, device registration and management, security, privacy and authentication, etc., are required. These tasks are accomplished by the so-called IoT Middleware (MW) that resides between the devices and the upper layers (e.g. cloud servers, consumers’ frontend), hiding the heterogeneity of the devices. Also, MW provides abstract layers for hiding the heterogeneity of the various medium access layers used (e.g. Bluetooth, LTE, ZigBee, etc.). A typical MW reference model is shown in Fig. 2.1.

Typical MW functionalities include code and service management, resource discovery and management, service integration, location tracking, semantic

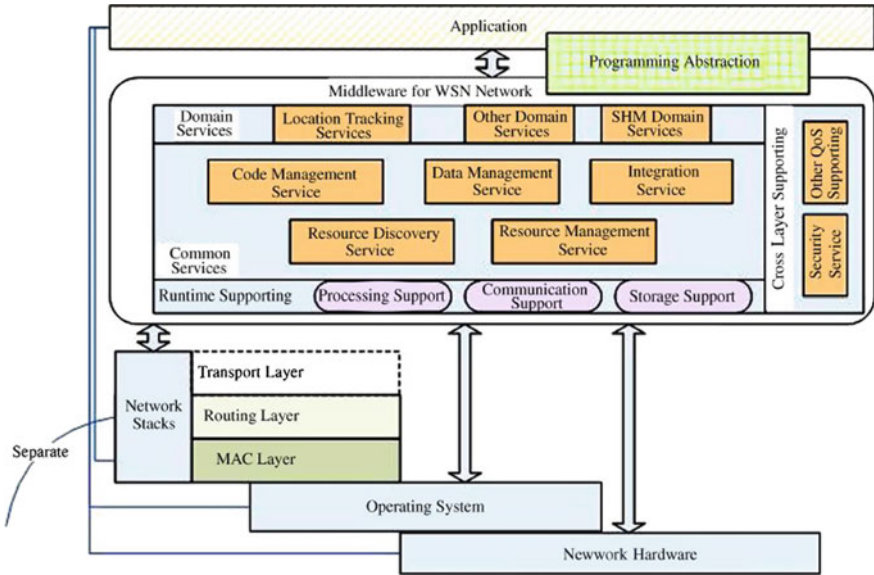


Fig. 2.1 IoT middleware reference model (Wang et al. 2008)

annotation, communication support, etc. Existing IoT MW architectures fall into three categories, namely the (Ngu et al. 2017): (i) *service-based*, where the service-oriented architecture is adopted, (ii) *cloud-based*, where the sensory data are stored in cloud servers and services like Testbed-as-a-Service, Measurements-as-a-Service, etc., are provided to users/consumers, and (iii) *actor-based* that focuses on open, plug-and-play IoT architectures.

There is a large number of significant research contributions on the IoT area, like OpenIoT (Soldatos et al. 2015), Node-Red (O’Leary and Conway-Jones 2017), Google Fit (Google Fit n.d.), Paraimpu (Pintus et al. 2012), etc. However significant challenges still remain. Key essential properties for a robust IoT MW are *security and privacy*. Due to the resource-constrained nature of sensors, the adoption of common cryptographic primitives and privacy-enhancing techniques is not always feasible. Strong encryption and privacy-enhancement is required for the transaction between the sensors and the MW. Furthermore, user authentication and authorization operations have to be supported by the MW.

2.2.4 Open Data Middleware Approaches in Deployed Smart City Infrastructures

Although IoT architectures are primarily evolving over network technologies and IoT middleware, currently the exponential growth of devices and sensors is leading to the generation of vast amounts of data from the smart infrastructure.

This inevitably has created the need for specialized data processes, such as storage, inference and analysis of the large amount of information. In this context, research is currently employing the power of semantic web technologies (e.g., SQL query (Lee 2010), SPARQL query (Apolinarski et al. 2014), RDF/RDFs and OWL languages (Zygiaris 2013) on top of linked data repositories. Such smart cities frameworks are discussed below.

Gambas (Apolinarski et al. 2014) is a middleware for the development of smart city applications that supports data acquisition, distribution and integration. The platform was deployed for several months in the transportation domain in Madrid, Spain. A Semantic Data Storage (SDS) component is used to store the data in the form of Linked Data. Query processors are used for data exposure to services and applications. The execution of these queries can be applied with the help of SPARQL query languages. In environments such as smartphones, the execution of query can be applied with the help of RDF-on-the-go (Le-Phuoc et al. 2010).

Sentilo (Bain 2014) is a platform for the management of sensors and actuators, deployed in the city of Barcelona, Spain. The platform includes many features, such as Cloud Computing, a non-SQL database, a memory database and a simple RESTful interface. Sentilo is designed and developed using open source components, such as Redis, MySQL and MongoDB databases, Hibernate, JSON or JQuery. A weak point of this platform is the lack of real-time data analyzing, obtained from different sensors.

Anthopoulos proposes an architecture that identifies a model for urban information, deployed in Kyoto, Amsterdam, Copenhagen and Trikala (Anthopoulos and Fitsilis 2010). This system follows a Service Oriented Architecture (SOA), where all services are stored and presented to stakeholders. The proposed architecture consists of five layers: the infrastructure, information, service, business, and stakeholder layers.

Zygiaris (2013) proposes an architecture that can be used in smart urban planning. This model was deployed in Barcelona, Amsterdam and Edinburgh. An important aspect is the storage and access of applications using Cloud Computing technologies. Semantic Web services and ontologies can provide an important interoperable data representation standard, while languages such as RDF-S and OWL allow the exchange of data across city domains. Data is exposed with the help of Visualization APIs.

WindyGrid (Rutkin 2014) is a platform for Smart Cities that presents real-time historical data, deployed in the City of Chicago. Specifically, the platform provides three main systems to the city of Chicago, which are: situational awareness and incident monitoring, historical data analysis and advanced real-time analytics. Big Data technologies were used for developing the platform, such as MongoDB, NoSQL database and parallel data processors. Examples of data include traffic conditions, buildings' information, and logs of emergency calls.

Table 2.1 contains an overview of the core functionalities and technologies of the aforementioned frameworks.

Table 2.1 Functionalities for smart cities' platforms

Smart city architectures/ platforms	Data acquisition	Data management	Data processing (service layer)	Data storage	Smart city deployed
Gambas	✓	✓	–	Semantic Data Storage (RDF, RDF-on-the-go)	Madrid
Sentilo	✓	✓	–	MySQL, MongoDB	Barcelona
Anthopoulos and Fitsilis (2010)	✓	✓	SOA	Mobile or Social network storage	Trikala, Kyoto, Amsterdam, base Copenhagen
Zygiaris (2013)	✓	✓	✓	Ontologies (RDF-s, OWL)	Barcelona, Amsterdam, Edinburgh
WindyGrid	–	✓	✓	MongoDB, NoSQL	Chicago

2.2.5 Data Analytics

IoT deployments for smart city applications are challenged by the need of dealing with the analysis of massive and heterogeneous data, in order the extraction of meaningful observations from raw sensing streams to be enabled. Towards this direction, addressing the accuracy of sensing data streams in real-time is of paramount importance for providing sophisticated services (Sun et al. 2016). This becomes even more evident as the community observes smart city platforms, being enriched with crowd-sensing models, wherein the use of error-prone, non-dedicated sensing elements are employed (Habibzadeh et al. 2017). Recent approaches in data analysis over IoT sensing streams are discussed in the rest of this section.

Considering the case of performing data analysis over a cloud-based architecture, Csáji et al. (2017) present a Smart City prototype, installed for monitoring pollution and traffic in Budapest. The respective architecture considers acquisition of information from the installed IoT sensing elements, and their integration over a software module, responsible for handling missing and noisy data; while being capable of creating short-term forecasts, accompanied by reliability estimates, in a batch-processing manner. Along the same direction, a four-tier architecture for smart city development and urban planning using Big-Data technologies (e.g., Spark Hadoop) is presented by Rathore et al. (2016). The objective of data management and analysis module is emphasized by combining historical information with real-time data for the prediction of future dynamic events. Similar principles, in terms of statistical data analysis, are also met in the City Data and Analytics Platform (CiDAP, SmartSantander project) (Cheng et al. 2015), which aims at processing both historical and real-time data.

Concisely, the CiDAP architecture defines a distributed computing framework and API for performing both internal (i.e., indexing, simple aggregation and first-order statistics) as well as external processing (i.e., clustering, anomaly detection, and classification). The information-based framework, presented by Jin et al. (2014) invests on sophisticated computational intelligence techniques (e.g., genetic algorithms, and neural networks) for converting information into knowledge, thereby silently implying that the raw streams of input data are a priori processed, cleaned, and normalized. Ultimately, the approach presented by Kolozali et al. (2014) elaborates on the reliable information processing over smart cities data, by means of patterns creation, fault tolerance mechanisms when malfunctioning or disappearing sensor are detected, and conflict resolution strategies when data analysis results in conflicting information; while it combines reliable (e.g., government) and non-reliable (e.g., crowd sourced) data.

Despite their scalable and modular architectures, the aforementioned approaches neither consider the inaccuracy of IoT urban measurements for the data analysis process, nor incorporate the extraction of on-line alerts while the system is in operation. In this sense, there exists a literature gap in extending statistical data analysis for IoT-based smart cities' applications beyond essential pre- and post-processing steps. As such, realistic factors that are sources of increasing level of uncertainty in raw IoT sensing streams, while enhancing the validity of detected alerting phenomena are not taken into account. Towards this direction, the herein employed statistical data analysis architecture differs from the current state of art in the following ways:

- It incorporates recent theoretical results that consider the characteristics of the sensing elements (e.g., accuracy, precision, resolution, sensitivity) for the lightweight and efficient quantification of the uncertainty;
- It adopts the respective toolbox as an inseparable part of the overall architecture for the on-line statistical data analysis;
- It employs the resulting uncertainty-aware information for the extraction of alerts that are associated to different types of system or data failure.

2.2.6 Information Visualization for Smart Cities

Information nowadays is rich and interconnected. It is actually not enclosed in a specific system, but is omnipresent in our surroundings, producing networks which create an Internet of Things (IOT). In this direction, the adoption of highly interactive visualizations holds a key role in providing insights on Big Data and IoT (LaValle et al. 2011). Both administrators and the public require intuitive visualizations, which provide access to data collected from IoT networks and illustrate information in an easily perceived manner. Application fields include energy management, networking, decision support systems, traffic monitoring and logistics (Singh et al. 2014).

2.2.6.1 OLAP—Charts

Big Data usually consists of multidimensional data sets, which are cumbersome to perceive and present. Online Analytical Processing (OLAP) is a widespread approach for interactively filtering out extrinsic information and its analysis, providing a clear view of the data from different perspectives. Pivot tables, also mentioned as cross-tabs, constitute a traditional interface for displaying OLAP data by employing a multidimensional spreadsheet in which a measure of interest is selected and corresponding dimensions present additional measures (Cuzzocrea and Mansmann 2009). Several types of plots are employed to illustrate multidimensional data, including numeric, ordinal, temporal and geographic values (Liu et al. 2013). Additional visualization methods include treemaps, circle packing, sunburst, parallel coordinates, streamgraphs and circular network diagrams (Wang et al. 2015).

2.2.6.2 Web Visualizations

Visualizations based on web technologies constitute an integral part of displaying cross-platform visual analytics in a manner familiar to both administrators and the public. Web analytics (Mikusz et al. 2015) are applied for gaining insights on statistical characteristics of IoT infrastructure and metrics. Additionally, web visualization is employed for spatial data representations. AVIoT (Jeong et al. 2015) is a proposed web interface that is suitable for both indoor and outdoor locations and can facilitate actuator placement in IoT environments. Visual analytics are proposed by Batty and Hudson-Smith (2014), combining chart representations, cartography, augmented and virtual reality in the context of urban design.

2.2.6.3 Geospatial Data

Geospatial data are present both in the context of Big Data and IoT technologies. The significance of spatial data characteristics was mentioned in 1970 by Tobler's first law of geography (Tobler and Waldo 1970), where "near things are more related than distant things", regardless of information interconnections. Geospatial big data constitute a significant portion of Big Data (Lee and Kang 2015).

Interactive maps are the prevalent interface applied for geospatial information, constituting a geographic visualization based on a common point of reference that displays information with regard to their location in space. Actions on the geospatial representations include map projection, pan and zoom (Cartwright et al. 2001). Common visualization approaches, utilizing map interfaces, include heat maps (Fisher 2007) and hypermaps (Kraak and Van Driel 1997). In terms of IoT visualizations, Merlino et al. (2014) combine maps and dashboards in order to visualize the sensory spatial distribution and the corresponding values respectively.

2.3 Heraklion Smart City Ecosystem

2.3.1 *Introducing the City of Heraklion*

Heraklion is the largest urban centre in Crete, the capital of the region and the economic centre of the island. The town enjoys a dynamic and imaginative combination of natural beauty climate, strategic position, cultural heritage and scientific background that has created a unique environment to support the broader entrepreneurial activity in the region and stimulate the local economy. Today Heraklion is the top choice for tourist destinations in the Mediterranean area, thanks to its strategic geopolitical position that connects three continents and many different cultures. According to the results of 2011 census, the population of the city was 173,993 inhabitants; while the Heraklion urban area has a population of 225,574 and it extends over an area of 225.5 km² (87 sq. mi).

According to a recent report published by travel analysts Euromonitor International (Euromonitor International 2017) and launched at the World Travel Market in London, Heraklion is Europe's fastest growing tourism destination for 2017. The city welcomed 11% more visitors in 2017, compared with the same period in 2016, surpassing its European rivals. In total 3.2 million visitors arrived in 2017.

The Mediterranean encompasses an amalgam of diversified nations and cultures. Despite this cultural barrier, there are similarities and common issues that can be found across the Mediterranean cities. Climate related issues, such as limited water supply and high temperatures during the summer, call for immediate action for an effective and sustainable way of water resource management. In addition, the current economic recession coupled with the refugee crisis across many Mediterranean countries requires new tools and methods for efficiently governing and managing cities and swiftly responding to an ever increasing demand for resources.

Challenges considered by this research work include the: (a) empowering of citizens by providing access to information and to new innovative services; (b) deployment of tools and services that will support intelligent decision-making for city management; (c) creation of added-value services for both citizens and visitors.

2.3.2 *Ecosystem Architecture*

Planning and developing a smart city ecosystem for the Heraklion case study required a careful and thorough examination of existing paradigms, architectures and technologies, such as the ones described in Sect. 2.2. Best practices in successful smart city instances were taken into account, leading to the development of

a modular platform that encompasses the integration of disparate systems and caters for future system scalability and expandability.

The Heraklion Smart City ecosystem (Fig. 2.2) consists of 3 *interoperable layers*, namely the Data Sources Layer, the Smart City Platform Layer and the Data Consumers Layer.

The *Data Sources Layer* contains the physical layer of the IoT infrastructure and the external Open Data sources, both of which act as data feeders of the Smart City Platform Layer.

The *Smart City Platform Layer* comprises:

- The core engine of the Middleware infrastructure, linking physical asset monitoring and external data extraction with databases and analytical engines.
- The Smart City web portal, serving public dissemination of accumulated knowledge.
- The monitoring and management services, providing smart city governance and city management tools.
- A public web API, available for third-party agents.

The *Data Consumers Layer* represents the Smart City end-users, namely:

- Public Administration
- Citizens and city visitors
- Commercial App providers

Access to these data represents an opportunity for SMEs to exploit and create added-value services for both citizens and visitors. A few examples of such services are travel applications based on real-time information, targeted advertising services

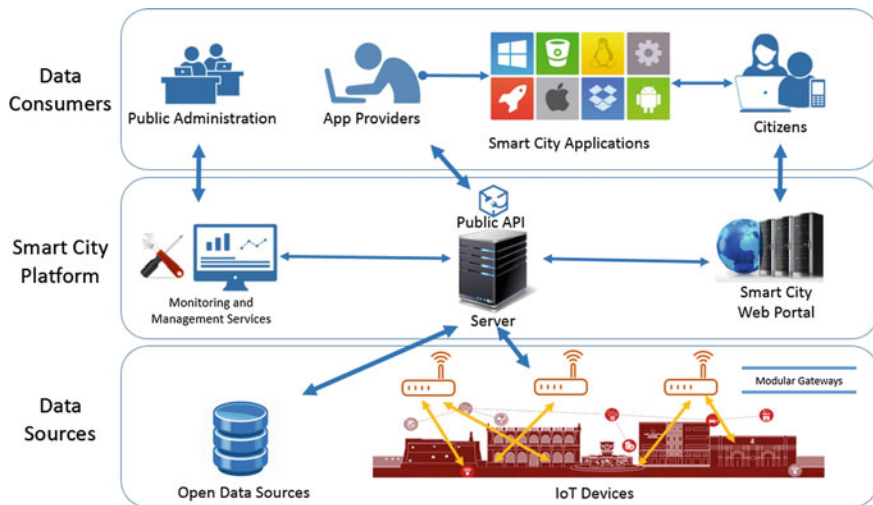


Fig. 2.2 Heraklion smart city ecosystem

and real estate tools that compare the potential value of alternative business or building locations. These data-enabled services could also provide a potential source of revenue for the data owners.

2.4 Heraklion IoT and Open Data Ecosystem Architecture

This section discusses the system architecture and the individual components that were integrated for building the core infrastructure of the Smart City ecosystem (Fig. 2.3).

2.4.1 Technology

The telecommunication standards, which may be incorporated in the case study architecture, are not fixed or predefined by the architecture. Any suitable communication technology can be integrated by the implementation of the proper software interfaces. Initially two standards were selected, based on their capabilities, performance, and robustness in large scale deployments, the LoRaWAN (LoRa Alliance n.d.) and the IEEE 802.15.4g (IEEE Std 802.15.4g-2012 n.d.), assisted by the 6LoWPAN (IETF n.d.) technology. In addition, this section covers the technological choices made for the platform's data repository and data analytics.

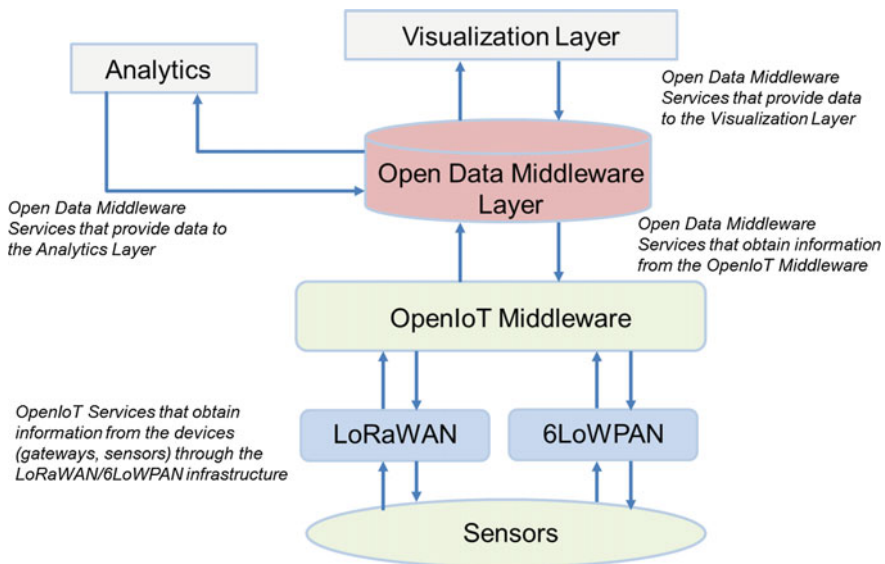


Fig. 2.3 Heraklion smart city architecture

2.4.1.1 Communication Standards Details

One of the most promising LowPower WAN technologies is LoRaWAN. It has been developed by a large multi-disciplinary consortium of companies; and it is based on the proprietary LoRa physical layer technology. It is based on the chirp spread spectrum (CSS) technology, providing many advantages over single carrier technologies, being very robust against interference.

LoRaWAN supports 6 spreading factors, from 7 to 12, depending on the selected data rate. The frequency bands for Europe are 863–870 MHz and 433 MHz, although there are no products yet to implement the 433 MHz band. EU868 MHz end-devices should be capable of operating in the 863–870 MHz frequency band and should feature a channel data structure capable of storing the parameters of at least 16 channels.

LoRaWAN networks are typically deployed in a star-of-stars topology, where Gateways, a.k.a. Concentrators, relay messages between end-devices and a central Network Server, which forwards the packets to the appropriate Application Server via the Internet. End-devices use single hop LoRa communication to one or more Gateways.

6LoWPAN is the acronym of IPv6 over Low-Power Wireless Personal Area Networks. It is an open IETF specification for the use of IPv6 networking over IEEE 802.15.4 based networks. Mesh routing, header compression and encapsulation are provided and optimized for the use with the 802.15.4 technology. Due to the use of IP technology, the connectivity of such networks to the Internet is seamless.

The set of standards under the umbrella of IEEE 802.15.4 is large. The one that currently provides a good balance between range and data rate is the 802.15.4g or Wi-SUN. It is optimized for very large scale applications use, mainly targeting on low power consumption, low data rate smart metering systems and advanced utility management systems.

IEEE 802.15.4g supports multi-rate and multi-regional frequency shift keying (MR-FSK); multi-rate and multi-regional orthogonal frequency division multiplexing (MR-OFDM); and multi-rate and multi-regional offset quadrature phase-shift keying (MR-O-QPSK) modulations. The frequency bands, which may be used in Europe, are 863–870 MHz and 2400–2483.5 MHz. The rate with respect to the modulation can be from 6.25 to 200 kbit/s.

2.4.1.2 Data Storage

The data storage component used by the Open Data Middleware (ODM) uses the relational database PostgreSQL (PostgreSQL n.d.) for storing the data produced by the IoT Middleware. PostgreSQL was selected over other database solutions for a number of reasons. It is a robust open source, object-relational database, with a very large community of users and developers. PostgreSQL carries many advantages such as flexibility, stability, usage of a standard data access language (SQL), support of ACID transactional consistency, limitless indexing, built-in data integrity

and a vast eco-system. Moreover, it supports a large number of advanced data types, such as multi-dimensional arrays, user-defined types, as well as geographic data types, provided by the spatial database extender PostGIS(PostGIS n.d.). PostGIS provides support for location SQL queries, a quite useful feature for IoT infrastructures, employed in the context of smart cities. PostgREST (PostgREST n. d.) is used for data exchange interoperability with the visualization layer and other 3rd party data consumer applications. PostgREST is a standalone web server that turns PostgreSQL database into a high performance RESTful API. PostgREST handles authentication via JSON Web Tokens (JSON Web Tokens n.d.); and delegates authorization according to the access control policies as specified and stored in the database, resulting in a single declarative source of truth for security.

2.4.1.3 Data Analysis

Going well beyond the calculation of first order statistics, the objective of the statistical data analysis herein employed is to timely detect abnormal changes in the sensing data streams, and enable early prediction of alerting phenomena. To this end, the High-level Data Management Toolbox (HDMA) (Tzagkarakis et al. 2014 and Tzagkarakis et al. 2015) is employed, originally designed and developed to meet the objectives of industrial Cyber-Physical Systems in general, and smart water networks in particular.

The key elements of HDMA (Fig. 2.4) are the quantification of uncertainties in given raw sensing streams; and the calculation of extreme events, considering the uncertainty-aware data. Concisely, the uncertainty in a quantity to be measured, henceforth called sensing modality (e.g., temperature), characterizes the dispersion of the values that could be attributed to that quantity (Aggarwal 2010). Thus, the uncertainty provides an indication on how inaccurate or incorrect a measurement is, as a result of the imperfections in the underlying sensing infrastructure.

The HDMA toolbox differentiates the independent sources of uncertainty into either statistical or systematic. The resulting values of uncertainty are consolidated

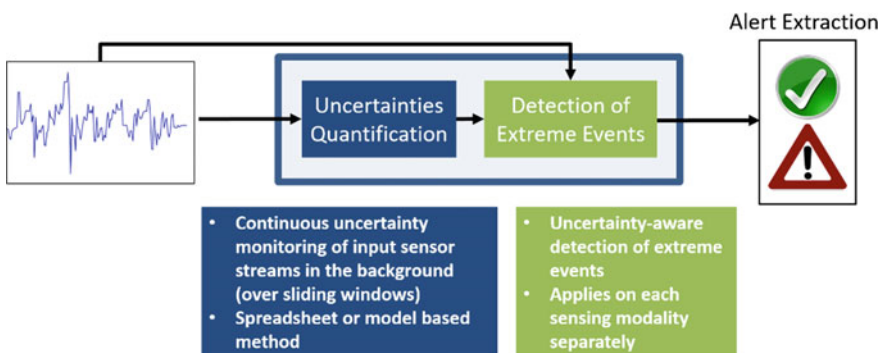


Fig. 2.4 The architecture of the HDMA toolbox (Tzagkarakis et al. 2014, 2015)

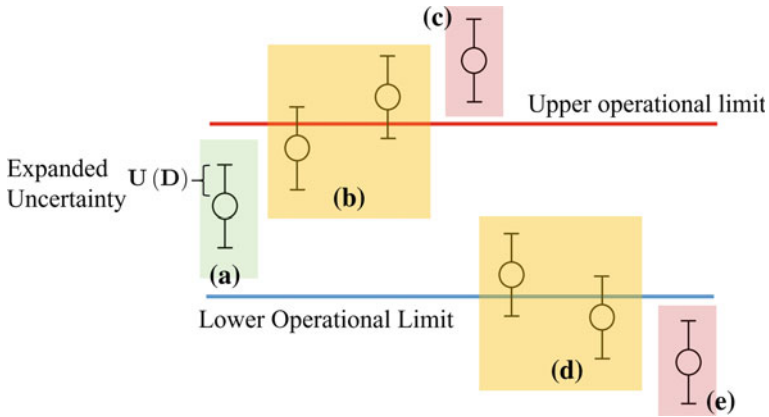


Fig. 2.5 The compliance with operating limit over the uncertainty-aware data stream (Tzagkarakis et al. 2014 and Tzagkarakis et al. 2015), indicating: **a** the case of full compliance with operational limits, **b–d** the uncertainty-aware value is below or above the upper operational limit, but the limit remains within the uncertainty, **c–e** the uncertainty-aware value exceeds the upper or lower limit

in order to calculate both the combined uncertainty, which characterizes the sensing modality; and the expanded uncertainty, which expresses the combined uncertainty at a specific level of confidence.

The expanded uncertainty for each sensing modality is in turn employed for enhancing the detection of extreme events over the data streams. Specifically, driven by the demand of providing timely notifications of alerting phenomena, the HDMA toolbox considers a modified version of the compliance-with-operating-limits method that incorporates the estimated uncertainty into the streaming data for detecting when the measurements exceed the upper or lower application-specific operational limits (Fig. 2.5).

While this method extracts an inference for each individual measurement of the input stream, the early-warning mechanism of the HDMA toolbox relies on a majority rule, according to which an alert is generated if N consecutive uncertainty-aware measurements exceed the operational limits.

2.4.2 Backbone Components

2.4.2.1 IoT Nodes and Sensors

This section provides a brief description of the hardware and software components that constitute the Nodes of the Heraklion IoT ecosystem. The nodes are distinguished into 6LoWPAN-enabled and LoRaWAN-enabled, since they differ in terms of wireless technology and smart-city services they support.

6LoWPAN Nodes

Each 6LoWPAN-enabled Sensor Node (SiSN) is built upon a Zolertia RE-Mote platform that hosts an ARM M3-Cortex running at 32 MHz, 32 KB RAM and 512 MB Flash Memory; and offers dual radio operation both in ISM 2.4 GHz and ISM 863–950 MHz frequency band, under IEEE 802.15.4 standard. External sensors can be attached to the RE-Mote platform through a number of communication ports (I2C, SPI, 12-Bit ADC with configurable resolution). The interested reader is addressed to (Angelakis 2016) for a more detailed characterization of selected sensor hardware.

IPv6 connectivity, specifically tailored for the low-power and resource-constrained nature of the RE-Mote platform, is offered by the Contiki OS (Dunkels et al. 2004). On the application layer of the network stack lays the Constrained Application Protocol (CoAP) (Bormann et al. 2012). Measurements collected through sensor drivers are exposed as CoAP resources by a CoAP server running on the RE-Mote. The CoAP asynchronous notification mechanism (OBSERVE) is used for periodic sensor measurement collection over UDP transport with application layer reliable unicast. Apart from sensory measurements, self-monitoring resources that report network statistics, device hardware/software info and power consumption are also exposed.

6LoWPAN-enabled Sensor Nodes interact with the OpenIoT Middleware (OMW) through the 6LoWPAN IoT Gateway (SiGW), which has fog characteristics such as local data processing, storage and networking services (Charalampidis 2017). It is mainly responsible for providing network and application protocol translation (translate a CoAP/UDP packet to a HTTP/TCP packet and a 6LoWPAN (IPv6) to a standard Ethernet/WiFi packet (IPv4) and vice versa). Apart from that, the SiGW offers functionalities such as sensor registration, monitoring and management, measurement aggregation and forwarding. Essentially, it hosts two different network interfaces. On the one side, there is an IEEE 802.15.4 interface, offered by a RE-Mote (acting as the 6LoWPAN border router) that enables connectivity to the Sensor Nodes. On the other side, connectivity to the OMW is provided by common interfaces, i.e. either the Ethernet or WiFi interface of a Raspberry-Pi 3 running Raspbian OS.

Each SiGW installed in the Heraklion IoT ecosystem ensures that the registration of the SiSNs to the OMW is performed in an easy, transparent and adaptive way. In particular, a CoAP server, running at the SiGW, plays the role of a registrar that handles registration messages received from the SiSNs; stores necessary identity information in a local database; and forwards registration messages to the OMW. Moreover, the SiGW implements a mechanism for per SiSN measurement collection activation/de-activation. Finally, additional security and reliability enhancing techniques include connectivity between OMW and SiGW, realized through the use of a VPN connection as well as local logging and monitoring, which ensures that transmission of measurements to the OMW is not disrupted by i.e. a reset of the gateway or a reset of a sensor node.

LoRaWan nodes

LoRaWAN-enabled Sensor Nodes (LoSN) are commercial-off-the-shelf (COTS) devices by Libelium (n.d.). There are two distinct categories of nodes, parking sensors and water quality sensors. According to Polycarpou et al. (2013), the expanding use of IoT technologies enables the collection of parking availability information even from on-street parking spots. The parking sensors are based on the measurement of the magnetic field strength in 3 orthogonal axes. The sensors are installed on the road surface and, at a programmable interval, they provide a status for the occupation of the parking space, based on the field strength compared to a predefined threshold. The water quality sensors are installed in city's water reservoirs and provide measurements about: temperature, conductivity, pH, water height, and chloride, ammonium, nitrate, calcium ions. All sensors are using the Microchip RN2483 LoRa modules, which are certified to the LoRaWAN specification (LoRa Alliance n.d.).

The respective LoRaWAN Gateways (LoGW) are built around the iC880A—LoRaWAN concentrator module. Based on the Semtech SX1301, the iC880A module is capable of receiving concurrently 8 LoRa packets with different spreading factors and channels. The module is connected to a Raspberry-Pi 3 computing device, combined with LoRa Gateway open source software provided by Semtech. The LoGW is providing the connectivity to the LoRaWAN server.

2.4.2.2 OpenIoT Middleware

As earlier mentioned, the OpenIoT middleware is a fundamental component of the IoT platform. The middleware realizes a layer between the IoT Sensor Network and the various applications that need to use the facilities of this network; and it seamlessly interconnects the IoT Network, exposing it as a virtual representation to end users, while also providing extra functionalities for improved security and privacy.

The OpenIoT middleware consists of several core components, namely (Fig. 2.6): the LSM Server, the Scheduler Core Server, the RabbitMQ Server and the RDF Server. The LSM Server and the Scheduler Core Server are deployed in an application server. The RDF Server and RabbitMQ server are standalone applications, which run on the same workstation. A rich HTTP Restful API is also used for serving communication among the middleware components of the Heraklion Smart City platform.

The LSM Server component is responsible for managing the IoT Network. One of its main functionalities is to listen for registration requests from the IoT network components, while registering and granting them access to the Heraklion smart city ecosystem. The notion of Features of Interest (FOIs) is employed. FOIs are digital representations of a physical or geographical object and logically represent the entity, which a sensor or a group of them is able to observe or manipulate through

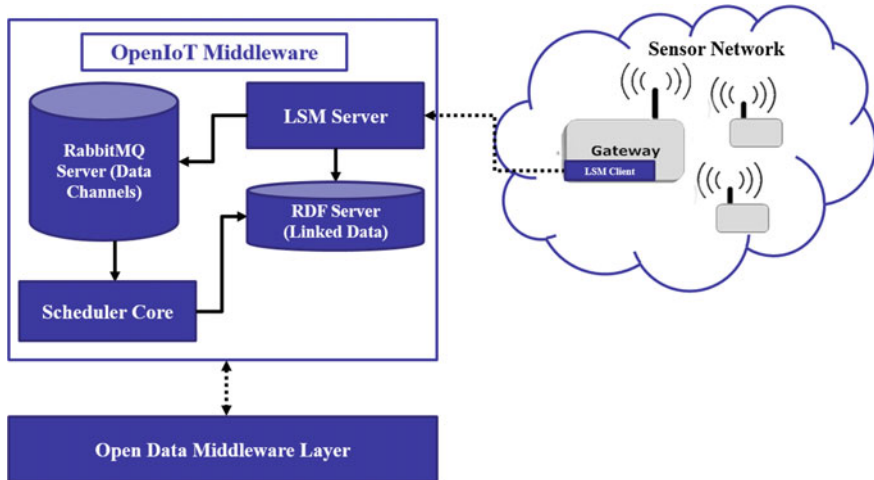


Fig. 2.6 Middleware components

sensing or actuation. The LSM Server registers IoT Gateways and IoT Sensors, while mapping them according to specific predefined FOIs. The LSM server component receives data from the sensors and publishes them to the RabbitMQ Server. It maintains the status of the IoT network and periodically updates metadata related to the network components.

The Scheduler Core component is responsible for handling the service requests from the Open Data Middleware (ODM) Layer. Due of the nature of the architecture, it provides an abstract view of the Testbed Topology to applications through random generated “handles”. This component also manages and maintains various data channels according to active handles and active data streams. The Scheduler Core HTTP Restful API provides functions, such as advertising a full list of available FOIs and the type of metrics that can be acquired from them. Other functions such as the ability to get a handle, based on a specific FOI, while getting data using it is also exposed. The Scheduler Core finds the mapped resource and creates the appropriate data stream from the resource that the ODM Layer has requested, thus generating and providing the handle on-the-fly.

The RDF Server is an instance of a Linked Data Server. OpenLink Virtuoso is being used, a scalable cross-platform server that combines Relational, Graph, and Document Data Management with Web, providing Resource Description Framework (RDF) capabilities. Its main purpose is to securely store and maintain data about the smart city topology (Sensors, Sensor types, IoT Gateways FOIs), such as geographical data, type of sensors per gateway, network information (IPs, ports), unique IDs for every component and their properties, etc.

For non-persistent storage, the RabbitMQ Server is used, which is an open source message broker software. Temporary buffers are created for those readings, which are relevant for an application or even multiple clients’ requests, and which

are meant only to store the sensors' data until their consumption. A time-to-live flag for each queue makes sure that data are not kept indefinitely. Using this technique, only requested data are maintained in the memory of the middleware, which auto-expire when they are not needed anymore, thus saving disk, memory and computing resources.

2.4.2.3 Open Data Middleware

The Open Data Middleware (ODM) is responsible for collecting, storing, maintaining and delivering data acquired from sensor data streams (through OpenIoT middleware) and external open data municipal sources (Fig. 2.7). Sensor data is retrieved by a set of Data Loaders, implemented in Java, separately for each modality (temperature, humidity, atmospheric pressure etc.).


Each data loader uses the OpenIoT middleware API to access and collect the recorded measurements, which, as described in Sect. 2.4.1.2, are then stored in a PostgreSQL database, a part of which is shown in Fig. 2.8.

Database entities are classified into the following four main categories:

Sensor network tables

These tables model the network of sensors that constitute the IoT infrastructure. Such tables are FOIS, DEVICES and SENSORS. The FOIS table holds all the features of interest (foi) that have been defined for the city of Heraklion, e.g. Eleftherias square, Morosini fountain square etc. Each foi can be linked with one or

Fig. 2.7 Middleware interoperability






Fig. 2.8 Open data middleware—database schema example

more DEVICES, which represent the physical devices in which different sensors are installed. The available devices and sensors, along with their characteristics, are stored in the DEVICES and SENSORS table respectively.

Sensor measurement tables

Sensor readings are stored in different tables, based on the modality of the measured value (e.g. Temperature, Noise, Wind etc.). Each modality table stores the time series retrieved values together with all related metadata (e.g. Fig. 2.8—AMBIENT_TEMPERATURE table).

Non-IoT Open Data tables

A separate set of tables is used for storing municipal open data from external sources. Such tables are shown in Fig. 2.9 and their content is described in details in Sect. 2.4.3.1.

Statistical analysis tables

This set of tables is reserved for storing the output of the Statistical Data Analysis Component. Attributes related to the analytics process, such as the temporal window of values examined, calculated uncertainties and identified alerts

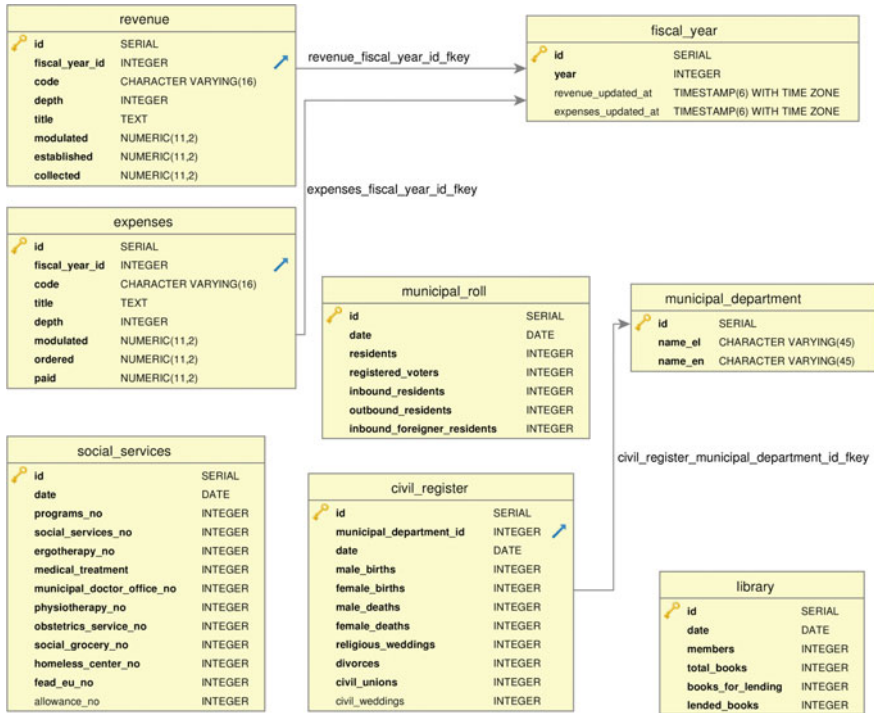


Fig. 2.9 Municipal data from external sources

(based on the Alert Level definition described in Sect. 2.4.2.4) are stored and maintained within these structures (e.g. Fig. 2.10).

As mentioned in Sect. 2.4.1.2, stored data is exposed to application level (both native and 3rd party applications) through a RESTful API defined by postgREST. The provision of complex aggregated data requests is accommodated by a set of user defined functions, aggregates and views. Examples are the:

- *avg_cyclic* user-defined aggregate function, which calculates the average wind direction.
- *get_wind_history* function that summarizes information about wind speed for a specific time period. Such function returns the occurrences of recorded measurements, for each secondary inter-cardinal wind direction (N, NNE, NE, etc.) and wind speed range in the Beaufort scale (calm, 1–2, 3–4, 5–6 and over 7). This function is used for visualizing wind historic data on a wind-compass schema.
- *dashboard_data* view returns the latest average values for all features of interest and for all modalities.
- *get_measurements_averages* function returns the average values given a modality, a set of devices, a time window and a granularity level of aggregation.

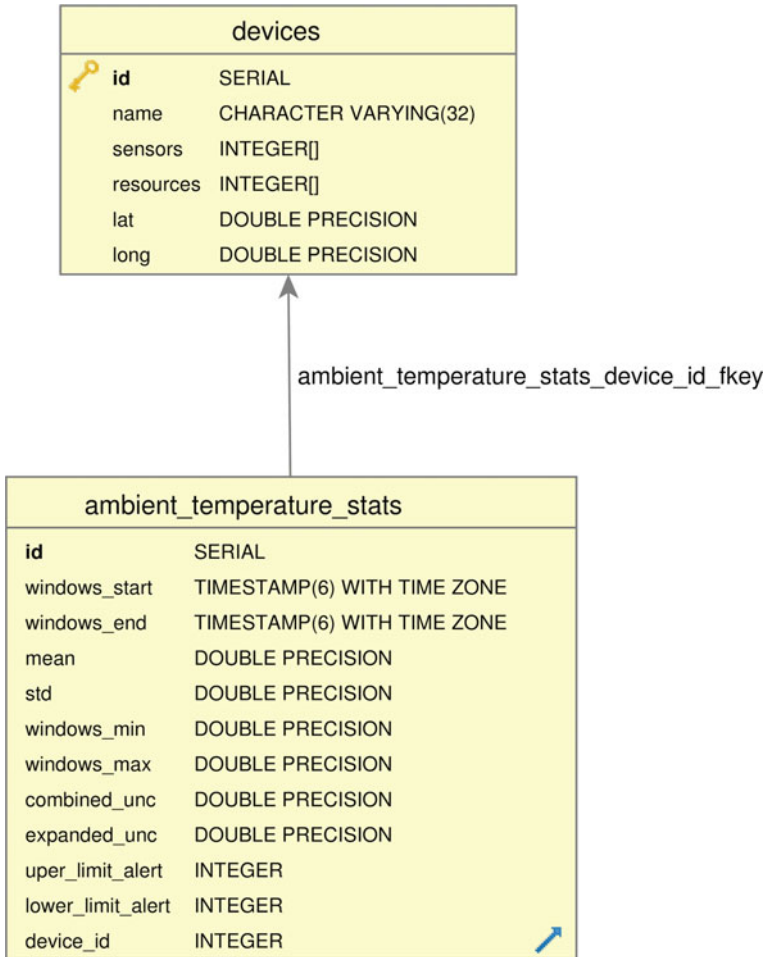


Fig. 2.10 Ambient temperature analytics

2.4.2.4 Analytics

A common concept for IoT deployments for smart cities’ applications is that sensing streams arrive at a centralized data management entity at extremely frequent time intervals from multiple locations. The herein architecture for the Statistical Data Analysis Component, presented in Fig. 2.11, considers this rationale for performing on-line quantification of uncertainties; and generating different levels of alerts.

Specifically, the Acquisition module is responsible for directly interacting with the Open Data Middleware and retrieving the sensing modalities $S_p = \{s_1, s_2, \dots, s_N\}$

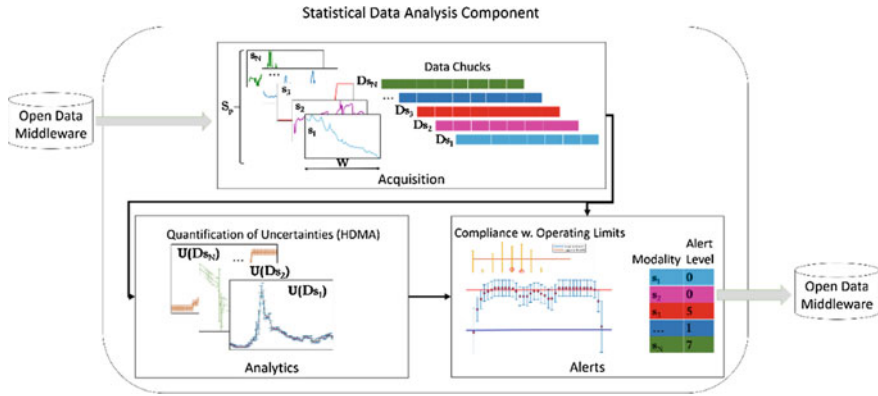


Fig. 2.11 The architecture of the statistical data analysis component

of each IoT sensing platform p deployed within the city. Depending on both the sampling rate of the sensing modalities and the on-line requirements of the end-user, these streams are derived in temporal windows, corresponding to W units of time (e.g., minutes). The resulting data chunk D_{s_j} for each sensing modality $s_j \in S_p$, is fed into the Analytics module, which is responsible for checking the contents of D_{s_j} for consistency, estimating the expanded uncertainty $U(D_{s_j})$ (Sect. 2.4.1.3), and providing first order-statistics for D_{s_j} over the temporal window W . Subsequently, the result of the Analytics module is employed by the Alerts module, which is responsible for categorizing the status the s_j -th modality of the p -th platform in different levels of alert, encoding different states of the s_j -th sensing modality:

- (a) The absence of data from the s_j -th modality at the specific temporal window W ;
- (b) The provision of a data chunk D_{s_j} with no statistical variance [$\text{var}(D_{s_j}) = 0$];
- (c) Non-compliance with the manufacturer operational limits;
- (d) Non-compliance with the application-defined limits for the uncertainty-aware data chunk $[D_{s_j} \pm U(D_{s_j})]$ (Sect. 2.4.1.3), which are associated to the application-defined limits (e.g., nominal range of temperature, healthcare limits for emission of gasses in the atmosphere).

Table 2.2 summarizes the different types of alerts considered.

The output of the statistical data analysis component (i.e., 1-st order statistics, expanded uncertainty and level of alert) is stored back to the Open Data Middleware. The above procedure is repeated for as long as fresh data streams arrive at the ODM, since the phenomena associated with these applications alternate at a frequent pace.

Table 2.2 The different types of alert that the statistical data analysis component considers

Alert level	Description	Reasoning of alert
1	Lack of data from the s-th modality within W	Platform/Network failure
2	No statistical variance of D_s	Sensor failure
3	Contents of D_s are below the lower operational limit of the sensor manufacturer	
4	Contents of D_s are above the upper operational limit of the sensor manufacturer	
5	$D_s \pm U(D_s)$ are above the lower operational limit defined by the application	Ambient of pollution emerging phenomena
6	$D_s \pm U(D_s)$ are below the upper operational limit defined by the application	
7	$D_s \pm U(D_s)$ are below the lower operational limit defined by the application	
8	$D_s \pm U(D_s)$ are above the upper operational limit defined by the application	
0	Normal status	

2.4.3 Heraklion Smart City—Open Data Web Portal

The Open Data Web Portal is a key component of the municipality’s strategy towards building a smart and open city. It serves as an information hub as well as a repository for storing and distributing the municipality’s open data (Fig. 2.12).




Fig. 2.12 Heraklion smart city web portal—home page

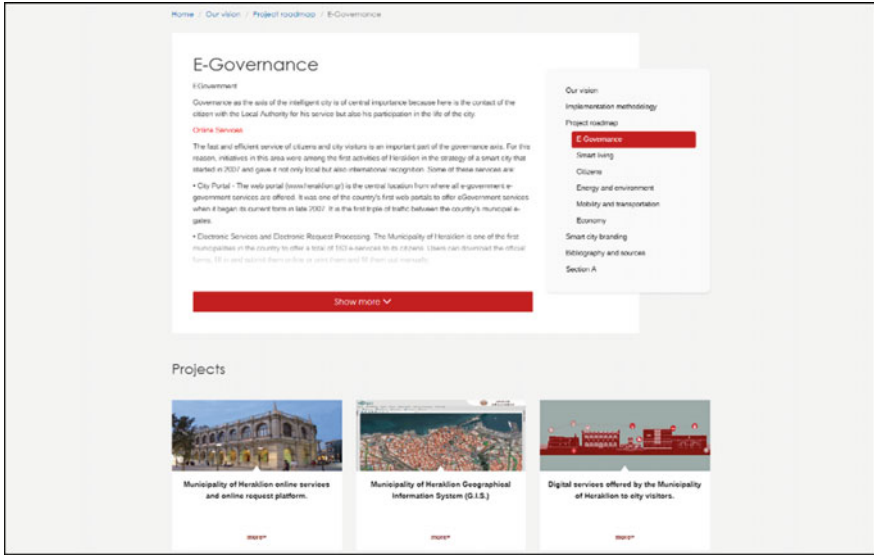


Fig. 2.13 Page with the detailed description of the e-Governance pillar together with the pillar’s related projects

The portal provides a detailed description of the municipality’s vision of a smart city. The vision aims to address the specific needs and challenges of the city; and focuses on six distinctive categories: e-governance, smart living, smart citizens, energy and environment, mobility and transportation, and economy. Citizens can view information regarding each category and related city projects that belong to the selected category (Fig. 2.13).

Additionally, the portal offers:

- News and updates regarding the municipality’s actions, events, festivals, etc.
- Updates about nominations and awards received by the city.
- Visualization of open data collected by the IoT infrastructure (either real time or by exploring historic data).
- Visualization of Municipal data.
- Access to Open Data Datasets.
- Web APIs for third party integration.

In addition, the portal provides a Management and Monitoring System (see Sect. 2.4.3.2), which on the one hand displays information regarding the current operational status of the IoT infrastructure (Gateways and individual devices’ status); and on the other hand, it contains an alarm notification system that informs the portal administrators of extreme events (e.g. when CO₂ health limits are exceeded) and of potential malfunction of specific sensors. The Management and Monitoring System can only be accessed by authorized users.

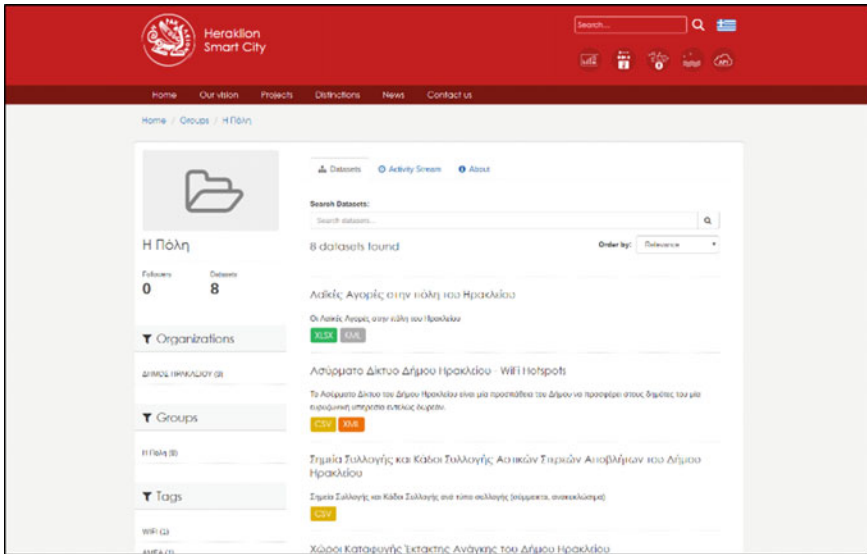


Fig. 2.14 CKAN dataset management system

On a technical level, the Heraklion Smart City Open Data Web Portal consists of three distinct components, the Content Management System, the Datasets Management System and the Open Data Visualizer. The three components are configured to provide a unified user experience, with seamless transitions between the components.

WordPress was selected as the Content Management System for its easy and comprehensive administration interface. CKAN was employed as the Datasets Management System (Fig. 2.14). CKAN is used by many public institutions seeking to share their data with the general public. The Open Data Visualizer was implemented as Single Page Application using Google's AngularJS framework. The graphs are dynamically generated SVG images, which are rendered using the D3js library.

2.4.3.1 Open Data Visualization and Exploratory Search

The Open Data Visualizer provides visualizations of sensor measurements, recorded by the IoT infrastructure and collected by the Open Data middleware, as well as data from other sources, such as the municipality's social service, the civil register, the demographics service etc. The Open Data Visualizer employs both infographics and basic charts (e.g. line charts, bar charts, etc.) for displaying information. Infographics are used as visual shorthand for presenting complex data quickly and clearly. Basic charts are used when the user needs an extended view of the collected

data. Basic Charts are presented alongside with exploratory search mechanisms for the examination and comparison of collected data.

The Visualizer employs the notion of Features of Interest (FOIs) for aggregating a set of devices and sensors under a single unit. Visualization of FOIs data is an aggregation of the recorded values of the individual devices or sensors in a specific geographic area.

The Open Data Visualizer is organized under the following categories: urban environment, water resources, parking monitoring and municipal open data.

Urban environment monitoring

The Urban environment monitoring component provides an interface for examining the Heraklion IoT infrastructure, displaying real-time data and exploring historic data. The following parameters are monitored: ambient temperature, ambient humidity, atmospheric pressure, wind (speed and direction), dust, VOC (volatile organic compounds), noise level, luminosity, and gases concentrations (SO₂, NO, NO₂, CO₂, O₃).

Water resources

Water sensors are installed on the municipality's water reservoirs. The water sensors record measurements regarding the quality and the quantity of the water (water level, water temperature, conductivity, PH and ions {NH₄⁺, NO₃⁻, Cl⁻, and Ca₂⁺}).

Parking monitoring

Parking sensors are placed in various locations in the city, where parking is prohibited. For that reason, the views for the parking sensors can only be accessed by municipal police. A hypermap is used to display the real time status of each parking sensor (either free or occupied). The map gets real-time updates by the Open Data Middleware for the current sensor status.

Other municipal data

Besides the data collected by the IoT infrastructure, the Open Data Visualizer also displays Municipal Data collected by other sources, such as the Civil Register, the municipal library, the social services, etc. Data is combined and presented in infographics that are generated dynamically by changing the selected time period (semester or year). Furthermore, individual dataset exploration is provided.

In order to enhance the user experience, several visualization components have been deployed with different user interfaces and functionalities, each serving a distinct scope.

Dashboard

The dashboard offers a quick overview of the real time data. The main goal of the dashboard is to display an integration of various data and to get the user familiarized with the different types of data that is being collected. The information displayed on the dashboard is organized in panels. Each panel is dedicated to a different aspect of the smart city ecosystem, such as the municipality, the citizens, the urban environment, etc. Users can interact with the panels and switch between



Fig. 2.15 The open data visualizer dashboard

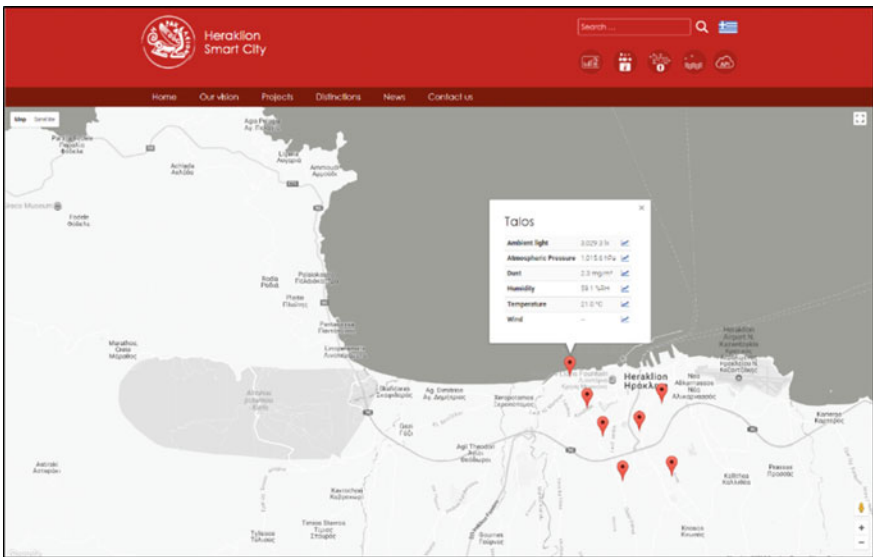


Fig. 2.16 FOI's with real time sensor measurements

views of the data being displayed. Each panel offers a link to a relevant page of the Visualizer, where the user can explore in detail the displayed data (Fig. 2.15).

IoT infrastructure

A hypermap is used for displaying the spatial distribution of the environmental IoT infrastructure. FOI's are displayed on a map that users can interact with, by panning and zooming as well as selecting specific FOIs. By selecting a FOI, a user can examine which parameters are being monitored and view their respective current values (Fig. 2.16).

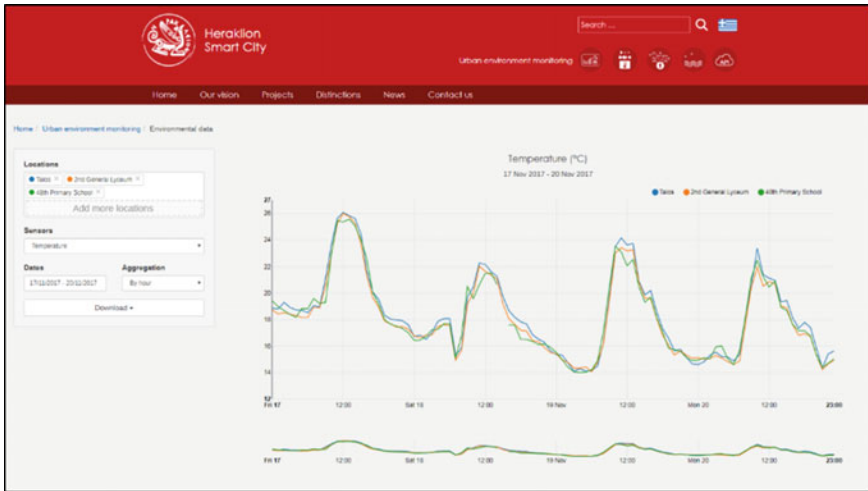


Fig. 2.17 The environmental data explorer

Live data

The Live data view visualizes the latest unprocessed (raw) values of the selected FOIs. Users can examine the real-time stream of data for all recorded parameters. Users can select the list of FOIs to be displayed; and adjust the length of the selected time window between one and six hours.

Environmental data explorer

Historic environmental data can be explored for each individual sensor modality. Data is displayed on a line chart with each line representing a FOI. The displayed values are aggregated by predefined time intervals. Users can adjust the selected time window as well the aggregation of the values. Possible aggregation values are by hour, day, week and month (Fig. 2.17). Moreover, wind measurements are additionally visualized using a compass rose (also known as wind-rose), which delivers a more concise perspective of the prevailing wind speeds and direction for a given location.

2.4.3.2 Managing and Monitoring System

Heterogeneous networks are systems encompassing components that differ in terms of software and hardware. A managing and monitoring system is required to present information regarding the status of the nodes, sensory data, network statistics, etc. The correlation of network statistics and weather condition can aid the network administrator to detect potential problems and take the necessary actions. A Monitoring Server (MS) has been developed for the Heraklion smart city ecosystem, collecting and displaying metrics acquired by the sensor network. MS is

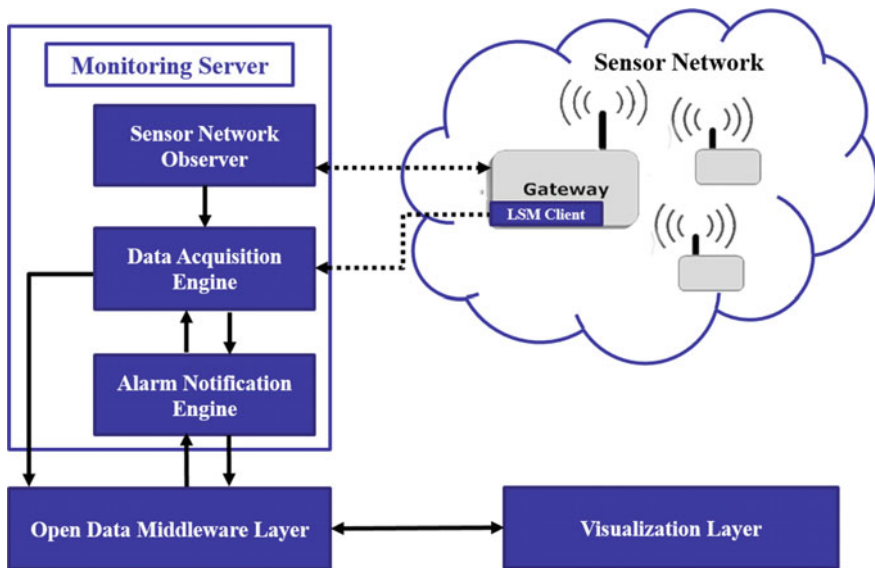


Fig. 2.18 Monitoring server architecture

also responsible for sending instant alerts in case of network failures or extreme environmental events.

The MS consists of the following modules (Fig. 2.18): The Sensor Network Observer, the Data Acquisition Engine and the Alarm Notification Engine. The goal is to develop a handy tool that incorporates all modules in a single instance; and can be deployed on various different platforms; hence it was implemented using NodeJS. The HTTP Restful API is used as the communication module of the system.

The Sensor Network Observer is responsible for acquiring data regarding the status of the Sensor Network. It polls the IoT Gateways periodically at configurable time intervals. At first, information such as the online status of the Gateway, the date that it was last seen online, the number of the respective sensors that it manages and which of them are online, are gathered by this component. The second level of information encapsulates data regarding the sensors, such as their unique ID, their Type and the types of measurements they are capable of providing, their IPv6 addresses, their status (online/offline) and the date their last activity was reported. HTTP GET requests are used for the polling procedure, and the format of information data is based on JSON.

The role of the Data Acquisition Engine module is to receive data from the Sensor Network Observer module, to aggregate it and transmit it securely in the Open Data Middleware Layer. In addition to that, it receives sensory data from the sensors, issuing HTTP posts. The IoT Gateways relay these measurements from the sensors to the Acquisition Module as soon as they receive them from the sensors they manage.

The Alarm Notification Engine is one of the basic parts of the Monitoring Server. Its main goal is to monitor the data, relying on a framework of rules; and alert the network administrators in case any rules have been violated. This framework is a logical combination of simple comparisons. For example, if the temperature is above 40 degrees or rainfall exceeds a certain threshold, then alert the administrator. Administrators are aided by the Alarm Notification Engine to take precautions or forecast an imminent failure of the network.

The extracted level of alerts for all sensing modalities over subsequent temporal window that are stored in the Open Data Middleware are fed into the Managing and Monitoring System of the Open Data Web Portal, for providing on-line notifications of the status of both the sensor modules as well as the ambient and pollution conditions around the city. Specifically, the Managing and Monitoring system yields notifications for the sensing modalities presented in Table 2.3. The lower and upper operational limits associated to either the manufacturer specifications or the nominal conditions¹ defined by the application (RERUM EU Project 2016) drive the extraction of alerts {3, 4} and {5, 6, 7, 8} respectively.

Table 2.3 The lower and upper operational limits for each sensing modality for the calculation of alerts

Sensing modality (unit)	Lower operational limit (manufacturer)	Upper operational limit (manufacturer)	Lower operational limit (application)	Upper operational limit (application)
Ambient light (Lux)	0.1	40,000	0	12,000
Loudness (dB)	34.6	95	34.6	85
Rainfall (mm)	0	10	0	8
Temperature (C)	-40	125	-10	50
Volatile organic compound (VOC) (ppb)	0	600	0	600
Sulphur dioxide (SO ₂) (ppb)	0	20,000	0	7.09
Ozone (O ₃) (ppb)	0	20,000	0	47.3
Nitrogen dioxide (NO ₂) (ppb)	0	20,000	0	19.7
Nitrogen oxide (NO) (ppb)	0	250,000	0	25,000
Particle matter (mg/m ³)	0	0.8	0	0.05
Humidity (%RH)	0	100	0	100
Carbon dioxide (CO ₂) (ppm)	0	20,000	0	430
Atmospheric pressure (hPa)	300	1100	600	1050
Wind speed (Km/h)	0	96	0	50
Wind direction (deg)	0	359	0	359

¹European Environmental Agency, <https://www.eea.europa.eu/>

Different levels of alert are addressed to different users of the Open Data Web Portal. Specifically, the levels of alert {1, 2, 3, 4} are associated to failures of the sensing or network underlying infrastructure (e.g., consistent packet losses, power/hardware failure) and are therefore aimed to the administrators of the platform. In contrast, the levels of alerts {5, 6, 7, 8} map to the emergence of meteorological (e.g., heat wave, flash flooding) or environmental (e.g., CO₂ emission) alerting, and thus are addressed to the environmental scientists and the municipal/prefectural civil protection agencies.

2.4.3.3 API—End User Documentation

The Open Data Web Portal exposes public web APIs available for consumption by third-party agents. The APIs endpoints provide direct access to data stored by the Open Data Middleware system and the Datasets Management system. The APIs are documented using the OpenAPI specification standard; and are offered with read-only access to the general public. The Swagger UI is used for rendering the API endpoints in the portal's frontend.

The Open Data Middleware API endpoints are grouped in the following categories: IoT Infrastructure, environmental measurements, environmental measurements analytics and other municipal data and datasets. The IoT infrastructure group provides endpoints, which deal with FOIs, devices, sensors, modalities and the relationships between them. The environmental measurements' group offers endpoints for accessing the raw, unprocessed data as recorded by the IoT infrastructure. The endpoints of other municipal data category deal with aggregated data regarding the demographics, the library, the municipality's financial data, the social services, etc. Finally, the dataset's group deals with retrieving datasets, dataset's files and their respective metadata.

2.5 Lessons Learnt—Future Work

Following the process of creating a Smart City ecosystem in Heraklion, several lessons were learned mainly stemming from the need to combine state of the art and non-mature technologies with out of the shelf-products. Especially in the case of IoT nodes, issues were faced with the extreme weather conditions of the Mediterranean (e.g. during summer), which resulted in the development of a new prototype node that could handle the large amount of sensors employed by the infrastructure, while at the same time being robust to weather conditions.

The locations where IoT devices are installed should also be very carefully planned, since the positioning of the devices can affect significantly the measurements, i.e. if the devices have direct contact with sunlight, the temperature or the light sensors may exhibit higher values. The installation points should also be selected in order to avoid physical tampering of the devices or excessive wireless

traffic that may create congestion and lost measurements. This latter part is extremely important, as most IoT devices operate in the ISM bands, which are normally congested within cities. This implies that communication links should be carefully tested prior to installation of devices at the desired locations. A reliable network monitoring mechanism should be in place so that the network administrator can easily check for malfunctioning devices or faulty links and act to solve the issues.

The configuration of IoT devices is mainly done manually with physical access to the devices (i.e. via a USB interface). Considering that in cities, hundreds or thousands of devices will be installed and many times in remote, not easily accessible, locations, it is not possible to easily reconfigure the devices manually. Thus, remote re-configuration is mandatory for IoT smart city installations and this can help significantly in the scalability of the overall system. Re-configurability can be used for solving bugs, changing security keys, adding drivers and supporting new protocols. This latter part is very important due to the severe technology fragmentation and the plethora of available communication and networking technologies and protocols.

This inevitably leads to the need for standardization, as existing IoT solutions follow different communication standards and platforms. Furthermore, as the infrastructure expanded, the need of data and device management became apparent. This can be perceived in terms of devices as the need for automated reporting, self-analysis, self-validation, self-operational assessment and ultimately self-healing (coping with the lack of human-resources in municipalities as a result of the economic recession in Greece). At the same time, the need for reporting was addressed through both the web portal and the VR visualization framework.

In the future, the ultimate goal is to disseminate the outcomes of this research work, while continuously updating the infrastructure to keep up with new scientific and technical developments. Updates in the IoT infrastructure include the deployment of more smart parking sensors throughout the city, city center traffic monitoring, municipal buildings energy management, etc. Furthermore, fostering the exploitation of the smart city ecosystem by the ICT sector at the city of Heraklion will create innovation and new added value services for all citizens.

References

- Aggarwal, C. C. (2010). *Managing and mining uncertain data*. Springer Science & Business Media.
- Angelakis, V., et al. (2016). *Analysis and evaluation of system performance and scalability*. RERUM Project, Deliverable D4.3, FP7-ICT-609094, European Commission.
- Anthopoulos, L., & Fitsilis, P. (2010). From digital to ubiquitous cities: defining a common architecture for urban development. In *Proceedings of 6th International Conference on Intelligent Environments*. <https://doi.org/10.1109/ie.2010.61>.
- Apolinarski, W., Iqbal, U., & Parreira, J. X. (2014). The GAMBAS middleware and SDK for smart city applications. *Proceedings of the IEEE International Conference on Pervasive Computing and Communication Workshops*. <https://doi.org/10.1109/percomw.2014.6815176>.

- Atzori, L., Iera, A., & Morabito, G. (2010). The internet of things: a survey. *Computer Networks*, 54(15), 2787–2805. <https://doi.org/10.1016/j.comnet.2010.05.010>.
- Bain, M. (2014). *Sentilo-sensor and actuator platform for smart cities*. Retrieved June 26, 2017, from <https://joinup.ec.europa.eu/document/sentilo-sensor-and-actuator-platform-smart-cities>.
- Batty, M., & Hudson-Smith, A. (2014). Visual analytics for urban design. *Urban Design*, 132.
- Bormann, C., Castellani, A. P., & Shelby, Z. (2012). Coap: An application protocol for billions of tiny internet nodes. *IEEE Internet Computing*, 16(2), 62–67.
- Caragliu, A., Del Bo, C., & Nijkamp, P. (2011). Smart cities in Europe. *Journal of Urban Technology*, 18(2), 65–82.
- Carrez, F. et al. (2013). *Final architectural reference model for the IoT v3.0*. EC FP7 IoT-A Deliverable 1.5, European Commission.
- Cartwright, W., Crampton, J., Gartner, G., Miller, S., Mitchell, K., Siekierska, E., et al. (2001). Geospatial information visualization user interface issues. *Cartography and Geographic Information Science*, 28(1), 45–60.
- Charalampidis, P., Tragos, E., & Fragkiadakis, A. (2017). *A fog-enabled IoT platform for efficient management and data collection*. In *Proceedings of the 2017 IEEE 22nd International Workshop on Computer Aided Modeling and Design of Communication Links and Networks—CAMAD* (pp. 1–6). <https://doi.org/10.1109/camad.2017.8031527>.
- Cheng, B., Longo, S., Cirillo, F., Bauer, M., & Kovacs, E. (2015). Building a big data platform for smart cities: Experience and lessons from Santander. In *Proceedings of the 2015 IEEE International Congress on Big Data* (pp. 592–599).
- Chourabi, H., Nam, T., Walker, S., Gil-Garcia, J. R., Mellouli, S., Nahon, K., Pardo, T. A. & Scholl, H. J. (2012). Understanding smart cities: An integrative framework. In *Proceedings of the Annual Hawaii International Conference on System Sciences* (pp. 2289–2297). <https://doi.org/10.1109/hicss.2012.615>.
- Contiki: The open source OS for the internet of things (n.d.). Retrieved June 28, 2017, from <http://www.contiki-os.org/>.
- Csáji, B. C., Kemény, Z., Pedone, G., Kutí, A., & Váncza, J. (2017). Wireless multi-sensor networks for smart cities: A prototype system with statistical data analysis. *IEEE Sensors Journal*, 17(23), 7667–7676.
- Cuzzocrea, A., & Mansmann, S. (2009). OLAP visualization: models, issues, and techniques. In J. Wang (Ed.), *Encyclopedia of Data Warehousing and Mining* (pp. 1439–1446), Second Edition. ISBN:1605660108, IGI Global.
- Dunkels, A., Gronvall, B., & Voigt, T. (2004). Contiki—a lightweight and flexible operating system for tiny networked sensors. In *Proceedings of the 29th Annual IEEE International Conference on Local Computer Networks* (pp. 455–462).
- Euromonitor International (2017). Top 100 city destinations ranking 2017. Retrieved June 29, 2017, from <http://blog.euromonitor.com/2017/01/top-100-city-destination-ranking-2017.html>.
- Fisher, D. (2007). Hotmap: Looking at geographic attention. *IEEE Transactions on Visualization and Computer Graphics*, 13(6), 1184–1191.
- Google Fit (n.d.). <https://developers.google.com/fit/>. Accessed 30 June 2017.
- Habibzadeh, H., Qin, Z., Soyata, T., & Kantarci, B. (2017). Large scale distributed dedicated-and non-dedicated smart city sensing systems. *IEEE Sensors Journal*, 17(23), 7649–7658.
- Hernández-Muñoz, J. M., Vercher, J. B., Muñoz, L., Galache, J. A., Presser, M., Gómez, L. A. H., et al. (2011). Smart cities at the forefront of the future internet. In J. Domingue, et al. (Eds.), *Future internet assembly* (pp. 447–462). Berlin, Heidelberg: Springer.
- IEEE Std 802.15.4 g-2012 (Amendment to IEEE Std 802.15.4-2011)—IEEE Standard for Local and metropolitan area networks—Part 15.4: Low-Rate Wire-less Personal Area Networks (LR-WPANs) Amendment 3: Physical Layer (PHY) Specifications for Low-Data-Rate, Wireless, Smart Metering Utility Networks.
- IETF (n.d.) IPv6 over low power WPAN (6lowpan). Retrieved 2 July, 2017, from <https://datatracker.ietf.org/wg/6lowpan/documents/>.
- INTER-IoT Project (n.d.). Retrieved 2 July, 2017, from <http://www.inter-iot-project.eu/>.

- Jeong, Y., Joo, H., Hong, G., Shin, D., & Lee, S. (2015). AVIoT: Web-based interactive authoring and visualization of indoor internet of things. *IEEE Transactions on Consumer Electronics*, 61(3), 295–301.
- Jin, J., Gubbi, J., Marusic, S., & Palaniswami, M. (2014). An information framework for creating a smart city through internet of things. *IEEE Internet of Things Journal*, 1(2), 112–121.
- JSON Web Tokens (n.d.). Retrieved 2 July, 2017, from <https://jwt.io/>.
- Kolozali, S., Bermudez-Edo, M., Puschmann, D., Ganz, F., & Barnaghi, P. (2014). A knowledge-based approach for real-time IoT data stream annotation and process. In *Proceedings of the 2017 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)* (pp. 215–222).
- Kraak, M. J., & Van Driel, R. (1997). Principles of hypermaps. *Computers & Geosciences*, 23(4), 457–464.
- Krco, S., Pokric, B., & Carrez, F. (2014). Designing IoT architecture (s): A European perspective. In *2014 IEEE World Forum on Internet of Things (WF-IoT)* (pp. 79–84).
- LaValle, S., Lesser, E., Shockley, R., Hopkins, M. S., & Kruschwitz, N. (2011). Big data, analytics and the path from insights to value. *MITSloan Management Review*, 52(2), 21.
- Lee, J. G., & Kang, M. (2015). Geospatial big data: challenges and opportunities. *Big Data Research*, 2(2), 74–81.
- Lee, Y. W., & Seungwoo R. (2010). U-city portal for smart ubiquitous middleware. In *Proceedings of the 12th International Conference on Advanced Communication Technology (ICACT)*, Vol. 1.
- Le-Phuoc, D., Parreira, J. X., Reynolds, V., & Hauswirth, M. (2010). RDF on the go: An RDF storage and query processor for mobile devices. In *Proceedings of the 2010 International Conference on Posters & Demonstrations Track*, Vol. 658 (pp. 149–152), [CEUR-WS.org](http://www.ceur-ws.org).
- Libelium (n.d.). <http://www.libelium.com>. Accessed 4 July 2017.
- Liu, Z., Jiang, B., & Heer, J. (2013). imMens: Real-time visual querying of big data. In B. Preim, P. Rheingans, & H. Theisel (Eds.), *Computer Graphics Forum* (Vol. 32, No. 3, pp. 421–430), Blackwell Publishing Ltd.
- LoRa Alliance, LoRaWAN 1.1 Specification (n.d.). Retrieved 2 July, 2017, from <https://www.lora-alliance.org/lorawan-for-developers>.
- Merlino, G., Bruneo, D., Distefano, S., Longo, F., & Puliafito, A. (2014). Stack4Things: integrating IoT with OpenStack in a smart city context. In *Proceedings of the 2014 International Conference on smart computing workshops (SMARTCOMP Workshops)* (pp. 21–28).
- Mikusz, M., Clinch, S., Jones, R., Harding, M., Winstanley, C., & Davies, N. (2015). Repurposing web analytics to support the IoT. *Computer*, 48(9), 42–49.
- Molinari, A., Maltese, V., Vaccari, L., Almi, A., & Bassi, E. (2014). *Big data and open data for a smart city*. Trento, Italy: IEEE-TN Smart cities white papers.
- Mulligan, C. E., & Olsson, M. (2013). Architectural implications of smart city business models: an evolutionary perspective. *IEEE Communications Magazine*, 51(6), 80–85.
- Natural Resources Defence Council. (n.d.) What are smarter cities? Retrieved July 2, 2017, from <http://smartercities.nrdc.org/about>.
- Ngu, A. H., Gutierrez, M., Metsis, V., Nepal, S., & Sheng, Q. Z. (2017). IoT middleware: A survey on issues and enabling technologies. *IEEE Internet of Things Journal*, 4(1), 1–20.
- O’Leary, N., & Conway-Jones, D. (2017). *Node red-a visual tool for wiring the internet of things*. Retrieved July 4, 2017, from <http://nodered.org>.
- Pintus, A., Carboni, D., & Piras, A. (2012). Paraimpu: a platform for a social web of things. *Proceedings of the 21st International Conference on World Wide Web* (pp. 401–404). ACM.
- Polycarpou, E., Lambrinos, L., & Protopapadakis, E. (2013). Smart parking solutions for urban areas. In *Proceedings of the 2013 IEEE 14th International Symposium and Workshops on World of Wireless, Mobile and Multimedia Networks (WoWMoM)* (pp. 1–6).
- PostGIS (n.d.). Retrieved July 5, 2017, from <http://postgis.net/>.
- PostgreSQL (n.d.). Retrieved July 4, 2017, from <https://www.postgresql.org>.

- PostgREST (n.d.). Retrieved July 4, 2017, from <https://postgrest.com>.
- Rathore, M. M., Ahmad, A., Paul, A., & Rho, S. (2016). Urban planning and building smart cities based on the internet of things using big data analytics. *Computer Networks, 101*, 63–80.
- RERUM EU Project (2016), *D4.3—Analysis and evaluation of system performance and scalability*. Retrieved July 4, 2017, from <http://cordis.europa.eu/docs/projects/cnect/4/609094/080/deliverables/001-RERUMdeliverableD43Ares20162302912.pdf>.
- RERUM: REliable, Resilient and secURE IoT for sMART city applications (n.d.). <https://ict-rerum.eu/>. Accessed 4 July 2017.
- Rutkin, A. (2014). How data can save a city. *New Scientist, 224*(2990), 24–25.
- Singh, D., Tripathi, G., & Jara, A. J. (2014). A survey of Internet-of-Things: Future vision, architecture, challenges and services. In *Proceedings of the 2014 IEEE World Forum on Internet of Things (WF-IoT)* (pp. 287–292).
- Soldatos, J., Kefalakis, N., Hauswirth, M., Serrano, M., Calbimonte, J. P., Riahi, M., & Skorin-Kapov, L. (2015). OpenIoT: Open source internet-of-things in the cloud. In I. Podnar Zarko, K. Pripuzic & M. Serrano, (Eds.), *Interoperability and Open-Source Solutions for the Internet of Things* (pp. 13–25). Springer, Cham.
- Sun, Y., Song, H., Jara, A. J., & Bie, R. (2016). Internet of things and big data analytics for smart and connected communities. *IEEE Access, 4*, 766–773.
- TinyOS (n.d.). Retrieved July 4, 2017, from http://tinyos.stanford.edu/tinyos-wiki/index.php/Main_Page.
- Tobler, W. R. (1970). A computer movie simulating urban growth in the Detroit region. *Economic Geography, 46*(sup1), 234–240.
- Tragos, E. Z., Bernabe, J. B., Staudemeyer, R. C., Ramos, J. L. H., Fragkiadakis, A., Skarmeta, A., Nati, M. & Gluhak, A. (2016) Trusted IoT in the complex landscape of governance, security, privacy, availability and safety. In O. Vermesan & P. Friess (Eds.), *Digitising the industry—Internet of Things connecting the physical, digital and virtual worlds* (pp. 185–214). River Publishers Series in Communications, ISBN: 9788793379817. <https://doi.org/10.13052/rp-9788793379824>.
- Tzagkarakis, G., Seliniotaki, A., Christophides, V., & Tsakalides, P. (2014). Uncertainty-aware sensor data management and early warning for monitoring industrial infrastructures. *International Journal of Monitoring and Surveillance Technologies Research (IJMSTR), 2*(4), 1–24.
- Tzagkarakis, G., Tsagkatakis, G., Alonso, D., Celada, E., Asensio, C., Panousopoulou, A., Tsakalides, P., & Beferull-Lozano, B. (2015). Signal and data processing techniques for industrial cyber-physical systems. In D. B. Rawat, J. Rodrigues, I. Stojmenovic (Eds.), *Cyber Physical Systems: From Theory to Practice* (pp. 181–226), CRC Press, Taylor & Francis Group, USA, ISBN 978–1-4822-6333-6. <https://doi.org/10.1201/b19290-12>.
- Wang, L., Wang, G., & Alexander, C. A. (2015). Big data and visualization: methods, challenges and technology progress. *Digital Technologies, 1*(1), 33–38.
- Wang, M. M., Cao, J. N., Li, J., & Dasi, S. K. (2008). Middleware for wireless sensor networks: A survey. *Journal of Computer Science and Technology, 23*(3), 305–326.
- Zanella, A., Vangelista, L. (2014). Internet of things for smart cities. *IEEE Internet of Things Journal, 1*(1), 22–32 (2014). <https://doi.org/10.1109/jiot.2014.2306328>.
- Zygiaris, S. (2013). Smart city reference model: Assisting planners to conceptualize the building of smart city innovation ecosystems. *Journal of the Knowledge Economy, 4*(2), 217–231.