

Smart Cities via Data Aggregation

Javier Poncela¹, Panagiotis Vlacheas², Raffaele Giaffreda³, Suparna De⁴, Massimo Vecchio³, Septimiu Nechifor⁵, Raquel Barco¹, Mari Carmen Aguayo-Torres¹, Vera Stavroulaki², Klaus Moessner⁴, Panagiotis Demestichas²

¹ *University of Malaga, ETSI Telecomunicacion, Malaga, 29071, Spain. {jponcela, rbarco, mdaguayo}@uma.es*

² *University of Piraeus, 80, Karaoli & Dimitriou Str, 18534, Piraeus, Greece. {panvlah, veras, pdemest}@unipi.gr*

³ *CREATE-NET, Via alla Cascata 56/D Povo - 38123 Trento – Italy. {massimo.vecchio, raffaele.giaffreda}@create-net.org*

⁴ *Centre for Communication Systems Research (CCSR), University of Surrey, Guildford GU2 7XH. United Kingdom. {S.De, K.Moessner}@surrey.ac.uk*

⁵ *Siemens Corporate Technology, Siemens SRL, Brasov, Eroilor 3a 500007, Romania, septimiu.nechifor@siemens.com*

Abstract Cities have an ever increasing wealth of sensing capabilities, recently including also Internet of Things (IoT) systems. However, to fully exploit such sensing capabilities with the aim of offering effective city-sensing-driven applications still presents certain obstacles. Indeed, at present, the main limitation in this respect consists of the vast majority of data sources being served on a “best effort” basis. To overcome this limitation, we propose a “resilient and adaptive IoT and social sensing platform”. Resilience guarantees the accurate, timely and dependable delivery of the complete/related data required by smart-city applications, while adaptability is introduced to ensure optimal handling of the changing requirements during application provision. The associated middleware consists of two main sets of functionalities: (a) Formulation of sensing requests: selection and discovery of the appropriate data sources; and (b) Establishment and control of the necessary resources (e.g., smart objects, networks, computing/storage points) on the delivery path from sensing devices to the requesting applications. The middleware has the intrinsic feature of producing sensing information at a certain level of detail (geographical scope/timeliness/accuracy/completeness/dependability) as requested by the applications in a given domain. The middleware is assessed and validated at a proof-of-concept level through innovative, dependable and real-time applications expected to be highly reproducible across different cities.

Keywords *resilient IoT, adaptive and scalable sensing, social IoT, smart city sensing toolkit, sensing zoom-in*

Corresponding author:

J. Poncela, jponcela@uma.es, ETSI Telecomunicacion, University of Malaga, Malaga, 29071, Spain. Phone: (+34) 952 13 27 59 , Fax: (+34) 952 13 20 27.

1. Introduction

Legacy cities are already characterized by the potential of collecting a wealth of data through a number of different sensing infrastructures. At present, this data is partially used for different, but often isolated, application purposes. Nevertheless, beside the bespoke applications and the associated sensing infrastructures that particular cities may have, the trend will continue to be towards an ever increasing amount of data coming from connected objects (sensors) as well as from “social sensing” applications. In other words, data are envisioned to come from (i) the Internet of Things (IoT) that becomes more and more established and (ii) the more and more sophisticated features of smart end-user devices in conjunction with the established social networks applications. Obviously, the effect of this trend will further enrich the overall “city sensing” capabilities.

In light of the above, the main question addressed in this paper relates to how to leverage such an increasing wealth of sensing capabilities to make an impact through sensing-driven applications (mainly from – but not limited to – “city management” stakeholders) that will effectively improve the quality of life of citizens in future smart cities. The main obstacle to overcome is that, at present, the vast majority of data sources are served through “best effort” data and data stream delivery, which is not enough for creating dependable, real-time and mission critical applications. Indeed, there is an increasing need for the applications to timely retrieve more complete and accurate data/data-streams, often in a dependable way. All these features, which are also in accordance with the Quality of Information (QoI) paradigm, can be seen as the fundamental issues to be tackled to enable an extended deployment of essential, innovative “city-sensing”-driven applications. The latter are envisioned to be, in turn, able to assist and support not only citizens’ everyday needs (e.g., real-time traffic sensing and environmental monitoring) but also those of city stakeholders (e.g., improving the management policies of the cities and providing prompt reactions against the identified problems).

The discussion so far yields the need to conceive a *resilient and adaptive IoT and social sensing platform*, whose associated middleware could actually play the role of key enabler to achieve the vision of better and more reactive services in a generic smart city framework of the future. We will see, in fact, that the proposed middleware is able to natively support the

selection of the most appropriate data sources and to be in charge of establishing and controlling the delivery of the data in agreement with configured constraints.

The rest of the paper is organised as follows. Section 2 presents the four basic pillars upon which the proposed middleware rests. Specifically, these are identified as a) virtualisation (that is, a dynamic and intelligent virtual representation of real objects and resources allowing them to become virtually “always-on”); b) cognitive management (that is, a dynamic selection of behaviour through self-management, taking into account policies, information and knowledge in the context of operation); c) resilience (that is, delivery of a service that can be justifiably trusted in dynamic environments) and d) scalability (that is, the capability to enable the coexistence of a large number of heterogeneous objects). Section 3 describes in detail the proposed middleware for a resilient, adaptive IoT and social sensing platform. Section 4 presents a set of applications that can be used to highlight the most representative features of the middleware. Finally, in Section 5 we draw some conclusions and outline possible future research actions.

2. Basic Pillars

This section discusses the four key themes that make up the enhancements of the proposed platform, namely: a) virtualisation; b) cognitive management; c) resilience; and d) scalability. Based on them, it aims to improve the representation and composition of IoT enabling technologies and to support and exploit a variety of enablers.

2.1 Virtualisation

A recent study on smart cities by the BBC points out that in 5 years’ time, the “*data infrastructure of our cities is going to become as important as the train and roads are now*” [1]. Recent advances in the wider IoT area also offer new ways to analyse and measure city life, from “*whether water pipes are leaking to how traffic is flowing on the roads and whether buildings are using energy in the most efficient way*” [2]. Examples include lamp-posts being fitted with sensors that can transmit information about cloud cover to offer hyper-local weather forecasting in Birmingham, to tweeting bus-stops in Norway to geo-tagging rubbish pieces to determine recycling efficiency in Seattle.

The above examples point to the ability of smart objects, such as sensors, to interact either among themselves or with other entities in the network. Facilitating such interactions requires provisioning mechanisms that enable virtualisation of such objects. Recommendations from

the IoT Special Interest Group of the Technology Strategy Board of the UK government also point to the need for design tools to abstract complexity and a layered, structured approach to data and resources [3]. Moreover, it is important to remember that the physical and digital aspects are considered as an integral unit in comprising a ‘thing’, though with various layers. In keeping with these observations, the virtualisation concept as applied in this paper envisages abstracting heterogeneous data sources into a common sensing substrate. The semantics used in the enrichment of virtualised object descriptions can facilitate knowledge extraction from available data, resulting in a substantially reduced amount of engineering needed for individuals or city administrations to provide new services. The envisioned services will leverage the increased wealth of city sensing data to provide the ability to easily adapt and control the required level of monitoring detail.

Existing platforms with connected devices and offering Application Programming Interfaces (APIs) for directory (searching objects) and data services include Xively (formerly Pachube/Cosm) [4], Nimbit [5], ThingSpeak [6] and Thingworx [7]. Research works include those applying semantic profiles to describe connected objects, as well as search strategies for the Web of Things (WoT) spaces [8] and frameworks following the Representational State Transfer (REST) principles to expose objects as web resources and gateways enabling such functionalities on devices running proprietary protocols, such as Zigbee [9, 10]. Several EU and international projects also exploit virtualisation concepts, such as CONVERGENCE [11] (concept of Versatile Digital Items (VDIs) as a common container for any kind of digital data, including representations of services, people and real world objects), EBBITS [12] (transforming every device into a web service with semantic resolution), SOFIA [13] (making physical world information available for smart services) and iCore [14] (virtually ‘always-on’ dynamic virtual representations of any real-world object) projects. The rise of citizen sensing is made possible by new social media that supports crowd sourcing approaches and when users converge in social media spaces in response to certain events [15]. Web-based social sensing platforms include the Ushahidi platform, which is a Google-maps based mashup for geo-located responses to a crisis and its instantiation in the Ushahidi Haiti map [16] which allowed people and organisations to post their needs on the crisis map following the 2010 Haiti earthquake. The WikiCrimes project [17] in Brazil aims to foster citizen and public organisation (e.g., non-governmental organisations (NGOs), newspapers) participation in law enforcement through crowd mapping. In the city context, the DiYSE project [18] has looked at deriving citizens’ behaviour patterns and obtaining feedback on a city-wide view of data

collected through Bluetooth-enabled City SensePod sensors which included noise, relative humidity, CO_x and NO_x levels and GPS sensors.

Virtualisation should include IoT sensing objects, either alone or in conjunction with associated external data stream processing, in a way that enables the generation of “smart” sensing feeds that can autonomously satisfy requests at different levels of sensing detail. It is worth mentioning that the sensing detail levels may reflect not only the geographical scope of analysis, but also other constraints such as accuracy, timeliness, etc. In addition to social sensing as enabled by the increasing number of features available on smart end-user devices, open data sources (e.g., those in the Linked Open Data cloud [19]) should also be seamlessly incorporated as sensing streams to enable smart cities. The virtualised objects need to be enriched with metadata as well as with cognitive capabilities, in order to achieve self-awareness, self-manageability and self-(re-)configurability. Representation of different levels of sensing abstraction and the ability to swiftly move between these (i.e., zoom-in and zoom-out feature) will also have to be provided in a domain agnostic way to empower similar applications in various smart-city applications.

2.2 Cognitive Management

The role of cognitive mechanisms has largely been applied so far to networks, with a cognitive network defined as a network with a cognitive process [20], also known as the MAPE-K cycle (Monitor-Analyse-Plan-Execute-Knowledge) [21], able to perceive current network conditions, plan, decide, act on those conditions and learn from the consequences of its actions. However, in the scope of IoT and in the context of smart cities, where potentially large sets of interconnected and heterogeneous devices produce data, appropriate cognitive techniques able to handle both devices and data in real-time and autonomously are more desirable. From theory to practice, the iCore framework and the associated platform [22] introduce cognitive management techniques in IoT at three levels, namely the Virtual Object (VO) level (self-configuration to maintain a constant link with the relevant real world objects and to manage the relevant data flows), the Composite Virtual Object (CVO) level (to enable the automatic selection of the most appropriate objects among the available ones for the resilient delivery of a service) and the Service level (to reason about the application requirements and the user preferences in a seamless way).

We believe that, to build reliable, smart and adaptive IoT applications for smart cities-, which needs to address the heterogeneity of a huge quantity of objects, as well as to guarantee the resilience of the associated sensing services in an automated way, the exploitation of

cognitive mechanisms is a prerequisite. Sensing devices should be intrinsically empowered with cognitive capabilities in order to achieve self-awareness, self-manageability, self-(re-)configurability during their operation. Thus today's sensing devices should be upgraded to future "smart objects". Different levels of sensing detail should be enabled to address the scalability in the data plane (being data processed as close as possible to their source), achieving not only the data sources selection and analysis according to geographical scope but also accuracy, timeliness, QoI objectives etc., as well as to reduce the overall amount of required sensed data to be delivered to the applications. In addition to smart sensing objects, cognition in terms of reasoning mechanisms is required in order to infer about the specific requirements/constraints imposed by the smart city application. These mechanisms will allow the automatic translation of the application into "machine/system parseable" requests for sensing services subject to constraints in terms of Service Level Agreements (SLAs) for timeliness and quality of information (e.g., accuracy) to be delivered. Learning mechanisms may also build knowledge on application preferences and eventually act on behalf of the application. Furthermore, optimisation/reasoning techniques are needed for the selection of the most relevant and trustful sources to be used for sensing data, as well as the more appropriate and resilient resources for the delivery of the sensed data to the application under the specified constraints in dynamic environments.

Citizens are also expected to contribute to smart city applications, thus making the step towards the "Social IoT (SIoT)", which will increase significantly the number of available sensing devices in a vast variety of geographical locations. This offers some level of redundancy to data being gathered, which gives an opportunity to assess both data integrity and take remedial actions that will complete missing information, adding robustness against failures and disconnections. To that respect, inference mechanisms are required for the retrieval of data from social sensing devices. Machine Learning (ML) and Complex Event Processing (CEP) can be used to aggregate and build knowledge from the captured data of sensors, breaking down "hierarchically" the problem of having to deal with a "flat" huge number of data sources. Semantic techniques may be exploited to select appropriate resources able to partially address application requests, when exact matches are unavailable. Learning techniques can help building knowledge about earlier decisions, so as to reach faster and/or more sophisticated decisions in similar future requests, as well as to acquire awareness about the available resources either in the IoT or in the Cloud infrastructure. Monitoring and assurance mechanisms for managing the real-time data flows are also of utmost importance, in order to either reason about a problem with a sensing feed or to predict probable anomalies,

as well as to provide remedial actions. Finally, intrinsic management functionalities, such as coordination techniques for handling the lifecycle of object virtualisations that support different applications but may control the same underlying resources, are highly required.

In a nutshell, cognitive techniques in the smart city context could be exploited to derive automated mechanisms for inferring about the application requests, for selecting and activating the most relevant resources for delivering sensing service at different levels of detail, as well as for managing the associated real-time data flows including Process Models to guarantee the availability and the integrity of data. The data sources considered include object virtualisations, cloud sources and social sensing streams. Last but not least, cognition should co-exist with security and privacy, allowing the owner of IoT sensing objects to decide to what “audience” and to which extent the owned/sensed data could be exposed.

2.3 Resilience

In its most general form, resilience is the desire to maintain dependability in spite of continuous changes. In our context it means guaranteeing the accurate/timely/dependable delivery of data. Social sensing may offer redundancy to sensing data and will add robustness against failures and disconnections.

Several reference models and architectures for self-adaptive systems have emerged in recent years. For instance, the multi-layer model proposed in [29] organises the system into a “*component control*” layer responsible for component-level adaptation, a “*change management*” layer for system-level adaptation based on a set of predefined plans, and a “*goal management*” layer for devising new adaptation plans when the predefined ones are no longer sufficient. The decentralised autonomic control patterns for self-adaptive software systems advocated in [30] aim at achieving quality requirements such as resilience, robustness and scalability in large distributed systems. The Quality of Service (*QoS*) *Management and Optimisation of Service-based systems* (QoS MOS) framework [31] introduces a reference architecture for self-adaptive service-based software systems. The QoS MOS architecture supports the dynamic selection of the concrete services that implement the operations of a service-based system to ensure that the performance and reliability QoS requirements of the system are achieved with minimal cost at all times. Finally, the authors of [32] identify a *unified information representation and exchange framework* as a fundamental principle of a cloud computing architecture capable of adaptation. The importance of this uniform information representation and message exchange between all business entities (i.e., cloud clients, partners and vendors) is further emphasized by the *information-centric* paradigm

defined in [33]. One approach could be introducing self-adaptive features at network and device levels. Another way is using an intelligent layer that, based on a set of rules, selects the best sources for each service according to user preferences.

In this area, new cognitive techniques to select and activate the most trustful and resilient sources are needed. Redundancy allows for robustness against failures and disconnections and provides the basis for remedial actions. In addition, anticipatory planning will help by pre-fetching and preparing of data in advance in fast changing environments. The research should also target new customized algorithms for scheduling, optimising and managing real-time data flows.

The approach presented in this paper intends to address resilience regarding both dependable services and real-time services that require low latency, as well as maximizing the availability of objects, the performance of services and the integrity of data. In this respect, the view pushes resilience one step forward by taking into account the QoI in addition to QoS and Quality of Experience (QoE). Therefore, in our view, resilience is the *ability to deliver service and information that can be justifiably trusted in spite of continuous changes and is in accordance with user requirements and resource constraints*. The latter reveals the need for mechanisms for adaptation to user requirements and resource constraints (e.g., time variant traffic demand, on demand provision and resource elasticity). As explained, the concept uses virtualisation and cognitive management as enablers to achieve resilience. The basic idea is to exploit the whole cognitive outfit as analysed for further enhancing the resilience capability. This includes optimisation/reasoning for the selection of the most appropriate sensing data objects (*use of trustful sources*) and infrastructure resources for the delivery of data under the specified constraints, situation awareness (in terms of objects, resources, application, etc.), cognitive monitoring and discovery, and triggering of appropriate remedial actions (assurance).

In this framework, to successfully enable resilience appropriate algorithms for scheduling, optimising and managing the real-time data flows should be investigated. Moreover, suitable Process Models should be better exploited as means to establish resilience (e.g., ask another source to send data, data caching, reestablishment of missing data, data reconciliation). A major contribution in this respect is the use of social sensing, which will significantly increase the number of available sensing devices across geographical locations. As stated in the previous section, this adds redundancy to the sensed data, which allows potential robustness and also the basis for supporting Process Models in the identification of remedial actions that will complete missing information. Besides integrity and reliability of sensed data, if these are

expected to support dependable applications, special care has to be taken also about the delivery of such data within given application constraints. Network virtualisation techniques (network slicing) are considered enablers, in order to exploit spare resources that may even belong to different networks, as well as to build a dynamic, federated virtual infrastructure to enhance the connectivity between the IoT sensing infrastructure and the applications that make use of it. Moreover, a flexible cloud infrastructure should be used to further support and enhance the IoT infrastructure (i.e., processing for data stream conversion) making it more robust against failure.

Another possible track is to investigate anticipatory planning for application requests through discovery of additional relevant objects that are deployed in the neighbourhood of the scope of interest and including them in a multi-plan. This feature can support the pre-fetching and setting up of data even before it is needed by means of real-time activation of these alternate data streams from the multi-plan by the middleware in sometimes fast changing environments (e.g., due to mobility, the source of data may become unavailable or just be useless), so that resilience (availability, integrity, performance) of data is guaranteed.

2.4 Scalability

Cities and urban areas represent the ideal testbed for future and innovative (or, in a word, *smart*) technologies, services and applications. However, if we consider that currently 50 percent of the world's population lives in cities (in Europe and the US, this percentage increases up to more than 75 percent) and that by 2050 this urban population will most likely reach 6.4 billion people [23], while already by 2020 we will reach the next trillion-level communication industry linking over a hundred billion devices [24], we understand that scalability represents a mandatory requirement for a future smart city context.

Loosely speaking, scalability means enabling the coexistence of a large number of heterogeneous objects (ecosystem). Thus, it leans on other basic pillars described in this section, namely virtualisation and resilience. Moreover, the virtual ecosystem will produce a huge amount of data, spanning environmental data produced by sensor networks, to video streams captured by traffic and surveillance cameras and by citizens with their personal smartphones. It follows that this data torrent has to be forwarded, stored and processed. This simple observation reveals interesting contact points with apparently orthogonal research and development planes, namely Social Sensing, Crowd Sourcing, Big Data and Cloud Computing, which are briefly summarized here.

Indeed, the widespread use of personal mobile communication devices (e.g., phones, PDAs) let people create and share content acquired by the on-board sensors (e.g., cameras, GPS) through the Internet, hence enabling new urban sensing and management opportunities, also known as Social Sensing [25]. Contextually, a completely new collective awareness is growing day by day and pushing people to collaborate towards continuous data collection, processing and interpretation, also known as Crowd Sourcing [26]. Thus, if judiciously analysed and aggregated, this huge volume of data of varying nature (e.g., raw data, text, image, video) coming from different sources (sensors, users of social platforms) at potentially high throughput can represent a gold mine of information to build effective smart city applications and services. It is not a coincidence that the three Vs that historically constitute the Big Data concept (Volume, Velocity and Variety) [27] were included in the previous sentence. Indeed, we believe that, to actually reach sufficient levels of scalability, smart city applications have to leverage on Big Data principles and tools. In particular, Big Data principles should be used to provide services at different scales (e.g., districts, towns, regions), while its data presentation tools (e.g., visualisation) could be used to guarantee added value for the citizens. Moreover, by means of Big Data tools, it is natural to think about individuals capable of aggregating open data with their own private ones, or defining their own services, composing existing ones, or even defining from scratch new ones. Finally, and as observed in [28], we believe it is a limitation to consider the Cloud as a mere endpoint of a smart city infrastructure. Rather, it should be handled in the same way as computing and storage resources usually are, in the more traditional Cloud stack e.g., abstracted, virtualized, and grouped into Clouds.

3. Middleware enabling a resilient, adaptive and social sensing IoT platform

The proposed middleware intends to support a wide range of “city-sensing” applications by providing them with resilient and adaptive data/data-stream delivery. In the previous section, we already introduced the notion of resilience, as it is conceived by the proposed middleware. As regards its ability to adapt, loosely speaking, it ensures the handling of new requirements that are likely to occur during a smart-city application provision.

The proposed middleware comprises the following two main sets of functionalities:

- Formulation of sensing requests and selection of the appropriate data sources. Data sources can be either IoT objects (simple ones or more intelligent ones as those

empowered by cognitive mechanisms (section 2.2)) or other devices, such as smartphones. This task involves the discovery of data sources, the assessment of the relationship of their data, and the realization of selection decisions, based also on the potential for the resolution of filtering and interoperability issues (the latter through associated processing at appropriate points on the delivery path).

- Establishment and control (activation, modification, release) of the necessary resources that are along the data path, according to the specific levels of timeliness, accuracy, completeness and dependability requested by the application. This task involves the activation of smart objects (i.e., virtualised sensors/IoT objects, jointly with the associated processing that adequately transforms the sensed data stream to fit the requests) and the selection of the suitable infrastructure resources (i.e., networks, computing and storage points) supporting the request.

The main proposition of the middleware is therefore an IoT sensing and resource control middleware or a “city-sensing” infrastructure. Figure 1 shows an example of how the proposed two functionality sets are supposed to interwork.

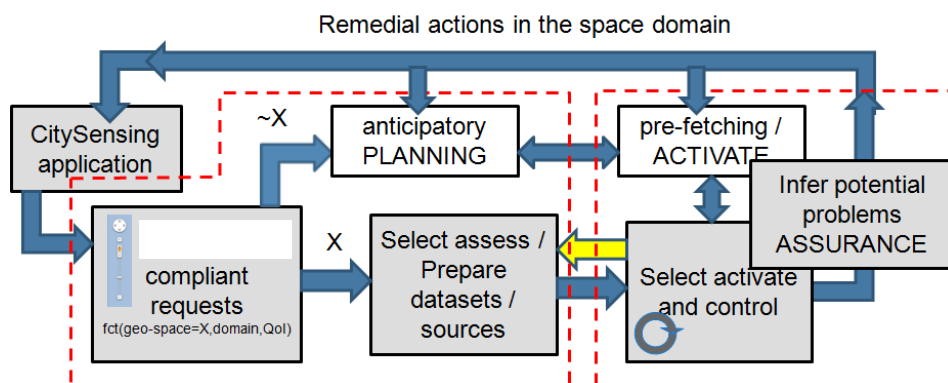


Fig. 1 Middleware high-level block diagram

City sensing applications submit their request to the interface component that parses the request. Such requests are formulated here in an abstract way as a function of X , representing the geographical scope of interest, the domain and the required quality levels. This request triggers the selection, assessment and preparation of the most suitable data sources; “suitability” here is not only meant to reflect geographical matching but also other constraints such as accuracy, completeness, timeliness, dependability, etc. If necessary, new smart objects are instantiated. The planning component integrates the most suitable data sources into a “multi-plan” that represents different alternatives for achieving the desired results. The

anticipatory planning component adds smart objects that may be needed if the city sensing application is interested in an adjacent area or in a different zoom level. The outcome is then passed for runtime execution and activation of resources on the delivery path. The runtime components use the plan to select and activate the smart objects that are most suitable, taking into account the current system situation and the optimisation criteria. Based on the anticipatory planning, the pre-fetching of information is initiated to enable a timely adaptation if the sensing application changes the request, as anticipated by the planning component.

The assurance component is there to “train” the middleware towards acquiring fast “zoom-in”, as well as “data repair” capabilities. Therefore, the results of the execution are returned to the requesting sensing application and, at the same time, feedback is provided to the anticipatory planning and pre-fetching components. The execution is being monitored and in case of problems, an attempt is made to find alternatives based on the existing multi-plan. Only if this fails, the more expensive re-planning is triggered by the planning component. In the picture “~X” is a representation of the neighbouring regions of X and the feedback arrows enable a quick reaction to zoom-in requests, as well as to requests for more accurate or complete data within a given region X.

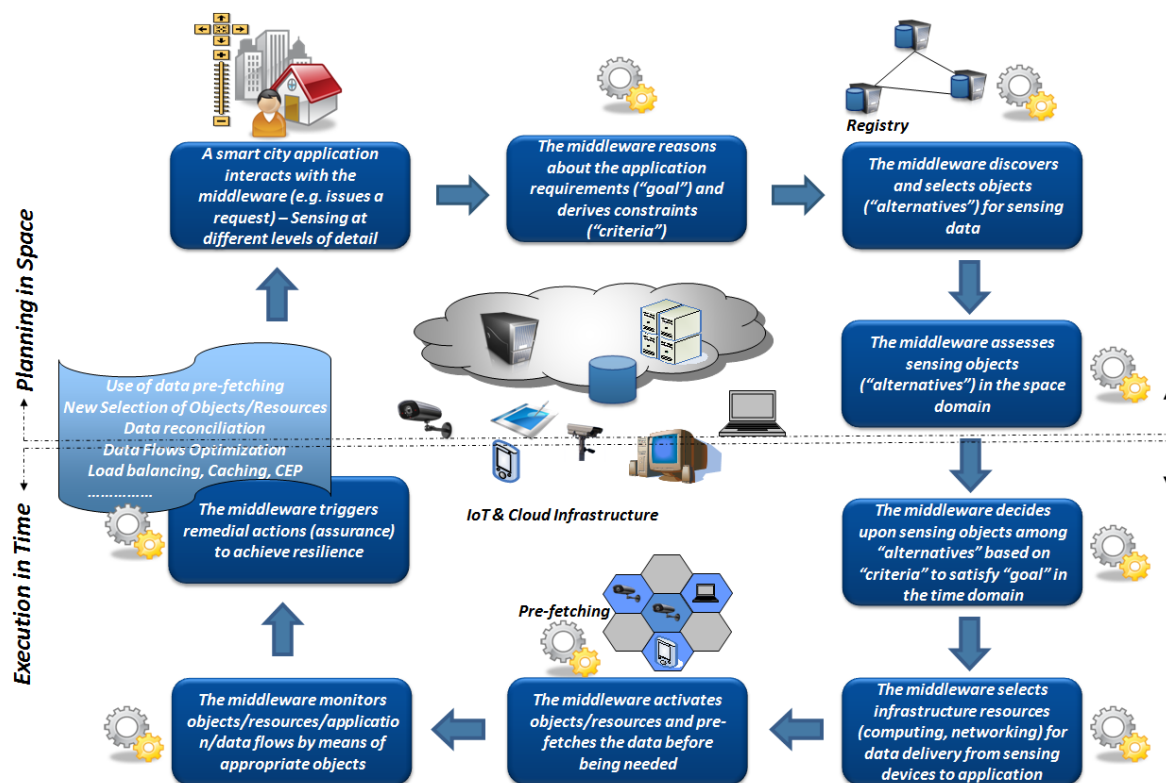


Fig. 2 Conceptual separation of middleware layers

Figure 2 provides another perspective on the functionality separation. There, the top half refers to the overall planning and preparation phase, responsible for the formulation of the sensing requests coming from various stakeholder domains. These requests involve subsequent discovery, inference and assessment of suitable data sources, as well as the type of processing needed to transform the associated data flows. The bottom-half of the figure refers to the overall execution and runtime phase. It depicts the establishment and control (activation, modification, release) of the necessary data flows, as well as the selection of appropriate supporting infrastructure resources (i.e., computing, storage and networking) for delivering the data to the applications within the given constraints of time, accuracy, completeness and dependability.

Special emphasis will be on developing a “level of sensing detail” associated with a zoom-in concept along different dimensions, able to trigger different outputs from IoT objects/sensors. These objects are virtualised through specific middleware features, making them more capable of supporting the production of different sensing output. The zoom-in concept and the anticipatory planning are key techniques for transforming an IoT based infrastructure into a flexible, efficient and scalable smart-city sensing substrate. Specially, the proposed middleware addresses scalability through the zoom-in feature that allows reducing the overall amount of required sensed data to be delivered to the applications, according to what they request. The use of cognitive techniques (ML and CEP) also addresses this problem by transforming the “flat” data source plane into a hierarchy.

Giving further insights into the zoom-in feature, let us consider the lowest level of detail. The target could be the rendering of an overall snapshot of “hot” areas on a city map (i.e., colour shaded mapping outlining areas with problems – high traffic density, abnormal energy/water consumption). Similarly the *middleware-enabled smart objects* could also be activated to produce detailed sensing information (i.e., “live video feed of traffic on a particular street”, or “current water consumption within a specific block of flats” etc.). A possible use of the proposed middleware would allow the selection of a region of interest on a map to produce sensing information in a given domain (i.e., traffic, energy, water) at a certain level of detail (see Figure 3).

In order to achieve these goals, it is envisaged that IoT objects will have to be virtualised in a way that exposes their features (including, when applicable, control interfaces appropriately secured for access) and either alone or in conjunction with associated external data stream processing, be able to generate *middleware-enabled smart* sensing feeds that satisfy requests at different levels of sensing detail. For example, a CCTV camera video stream can either be

The subsequent impact is assessed and validated at a proof-of-concept level through innovative, dependable and real-time applications (presented in the next section) that will help and support city authorities to better address citizens needs and requirements and which are expected to be highly reproducible across different cities.

4. Applications in real Smart City context

This section presents some smart city applications that demonstrate the usefulness and feasibility of the innovations proposed in this paper. Applications should take into account the cross-application nature of objects from several fields, such as smart living, smart spaces and smart communities. The aim of the applications in a smart city context could be to identify priority areas to improve quality of life, to detect and analyse public safety issues, to show city status with detailed analysis support, to introduce eco-friendly principles, to keep the highest quality standards for the tourism industry, to support the authorities in the analysis and solving of mid/long terms patterns of infrastructure failure, criminality, etc.

The applications make use of the technological infrastructure available in many cities: Wi-Fi networks, broadband network (optical fiber), cameras, Geographic Information System (GIS), public lighting network, traffic management system, parking sensors, water meters, etc. In addition, high amount of data flows from heterogeneous sources are involved and, from these flows, proper high-level views are obtained through the use of the smart processing algorithms proposed in this paper. The selected applications are depicted in Figure 4.




Fig. 4 Applications in real Smart City context

a) Smart and Green Mobility

The objective of this application is to achieve social and “green” mobility by organising traffic in a sustainable and rational way and by helping citizens optimise their time when moving within the city. With this aim, traffic information is collected from multiple heterogeneous sensing devices e.g., traffic sensors, traffic cameras, etc. Information from smartphones (SIoT), from electric cars and from public vehicles (e.g., buses, police, taxis) can be used in combination with other sensors’ information to reason about a traffic congestion. Parking sensors can be used to know the available parking places. A centralised system for incident management can be used to avoid crossing areas where incidents have happened or, conversely, to inform about incidents in the city. Traffic lights can be regulated based on traffic analysis, concentration of public transport, etc.

Once this information is processed and smartly analysed, maps of traffic intensity are derived and subsequently several applications are built, such as:

- 1) The user defines his final destination. Then, the application suggests the best route that he could follow to save time (e.g., avoiding traffic congestions) and guides him to a parking lot with available places in the area. Other constraints can be taken into account e.g., to follow a route with minimum air pollution.
- 2) The routes of the waste trucks can be optimally scheduled, taking into account traffic issues, the seamlessness of the collection in line with the social and financial fabric of the city/country, the resources consumed by the truck fleet, the data from the sensors in the waste containers, etc.
- 3) Intelligent management of public transport. Routes of buses can be optimised depending on traffic density (change routes, change bus frequencies, etc.). In addition, the taxis, or even buses, in a given area can be coordinated.

The innovations proposed in the previous sections are essential to efficiently develop these applications. Firstly, the huge amount of heterogeneous object instances (traffic sensors, traffic cameras, parking sensors, traffic light regulators, incident management system, sensors in waste containers, taxis, smartphones, credit card payment, etc.) involved in these applications require the capability of dealing with complex real-time data flows. Secondly, fast zoom-in both in time and space to highlight critical areas is needed. Thirdly, social sensing innovations are also an added value to these applications by using citizens’ smartphones as a rich source of data. Lastly, the timeliness and accuracy of the exchanged information is of critical importance, in order to realise resilience.

b) Smart and Green Management of Energy and Water resources

The objectives of this application are to manage the energy and water resources of a smart city in a smart and green way and to help the city and the citizens to reduce the energy and water billing expenses. With this aim, water and energy consumption should be measured in each house, something that is already under way. Different resolutions in time and space are defined for the measurements. For example, for billing purposes it is enough to take measurements once a month, but it has to be done in each and every home. Conversely, if there is an abnormal consumption in an area (e.g., due to a water leakage), measurements will have to focus on that area and need to have a higher time resolution.

Subsequently, maps of energy and water consumption in the city, together with indicators of water quality, are built and behaviour patterns are investigated. In addition, the citizens, by means of their IoT social network tools, are able to exchange information on their water and energy usage to perform local benchmarking, to warn about incidents, etc. The city could offer incentives to the citizens with bigger energy and water consumption, as well as encouraging participation.

The large amount of sensing devices and the huge amount of information to be remotely measured will need the features described in previous sections, in order to select the appropriate support infrastructure resources for the delivery of the sensing data within given constraints of time, accuracy and quality. In addition, these applications will make use of the proposed capability of zooming in those areas which have higher energy or water consumption, especially when there is an abnormal increase that can indicate an incident. Social sensing innovations will also be an added value to this application by using citizens' smartphones to exchange information about water and energy usage in those areas where remote monitoring is not possible.

c) Smart public safety

The objective of this application is to provide an effective way to detect, analyse and solve public safety issues (emergencies, criminals on the run, lost children or children in danger, outbreaks in public manifestations, hazardous material transports, etc.). This should be done by supporting citizens to provide real-time communication with the authorities about these issues, by informing citizens and evacuating/isolating the involved city areas and by supporting the authorities in the analysis and solving of mid/long-term patterns of infrastructure failure, criminality, etc.

Citizens can report to the authorities about public safety issues such as fires, water leakages, robbery or similar incidents through a centralised incident management system. With the enhancements proposed in this paper, the dynamic capabilities of citizens

offering “events” out of their own data stream captures (sound, movement of the smartphone, video, etc.) will be added. In addition, traffic and security cameras, dynamic user reports, and other data stream sources will allow the public authorities to detect, monitor and track the geographical evolution of the public safety issues. Some example areas where this could be applied are:

- 1) The density of people in an incident area could be estimated based on multiple sensing devices (cameras, electronic payment, entrance in museums, smartphones, etc.). After the intelligent aggregation of the resulting real time information, public authorities will be able to quickly determine the risk for the population, so as to decide in a well-informed way e.g., on the number of staff needed to control the area, on the amount of temporary sanitary installations needed, etc.
- 2) Support for the determination of a traffic-density-optimised path that the emergency response teams can follow to reach the incident area in the fastest way.
- 3) In the mid/long-term, the daily routes of the police patrol can be scheduled so as to optimally anticipate incidents in areas known to be most problematic. Based on the proposed features, these routes could be automatically redefined following the dynamic variations of identified criminality and conflict patterns.

These cases illustrate well the use of a number of the capabilities proposed in this paper. First of all, for example, in incident reporting and monitoring, a wide, heterogeneous range of device types and data sources, such as city or traffic cameras, smartphones, incident management system inputs, etc. needs to be supported. Secondly, for authorities to take informed decisions, real-time data analysis is needed, dynamically taking into account all relevant devices and data sources, but also with the system “understanding” potential needs of the authorities in this respect. The automated collection (and learning) of real-world knowledge, as part of the planning capabilities of our proposed solution, and the efficient aggregation and processing of all the needed data (based on dynamically deployed CEP technology) are key enablers needed for this. Thirdly, the proposed “zoom-in” functionality will enable to monitor a moving target of interest in an efficient way, with the allocation of processing resources dynamically following the important aspects, varying geographically or in data type nature, as required by a response team. The resilience offered by the features proposed in this paper is also important in order to deliver the service in a fast changing environment e.g., due to mobility of targets and sensing devices, as well as the real-time nature of the exchanged information. Finally, since it would be difficult to cover all the city area with sensing devices, the integration of

social sensing capabilities in the platform will allow increasing the number of sensing objects and using them in an opportunistic manner.

Conclusions

In this paper we presented a resilient, adaptive and social sensing platform for smart cities, in order to address the pressing need for on-time and dependable delivery of data/data-streams in innovative “city-sensing”-driven applications. This requirement is crucial for supporting not only citizens’ everyday needs but also the needs of the respective city stakeholders. The basic pillars of the associated middleware were identified in terms of virtualisation, cognitive management, resilience and scalability, together with the added value features that the pillars may offer to such a platform. The high level principles of the middleware were described in detail, including the discovery and selection of appropriate data sources, the establishment and control of the data delivery in agreement with configured application constraints, the increase of the number of available sensing sources through social sensing and the ability to produce sensing information at a certain level of detail (e.g., geographical scope, timeliness, accuracy, completeness, dependability) according to the application request. The middleware was validated at a proof-of-concept level, through the presentation of innovative, dependable and real-time applications, namely Smart and Green Mobility, Smart and Green Management of Energy and Water resources and Smart Public Safety, which are enabled by the middleware added value features. These applications make use of the technological infrastructure available in many cities, such as Wi-Fi networks, broadband network (optical fiber), cameras, Geographic Information System (GIS), public lighting network, traffic management system, parking sensors, water meters, etc., and they could exploit the smart processing algorithms to effectively maintain and handle the high amount of data flows produced from such numerous heterogeneous sources. Future plans include a more thoughtful investigation of security and privacy concerns, as well as a more detailed proof-of-concept in terms of qualitative/quantitative results.

References

- [1] J. Wakefield. (2013, 19 August). Tomorrow's cities: Do you want to live in a smart city? Available: <http://www.bbc.co.uk/news/technology-22538561>

- [2] J. Wakefield. (2013, 27 August). Tomorrow's cities: How big data is changing the world. Available: <http://www.bbc.co.uk/news/technology-23253949>
- [3] "Internet of Things (IoT) and Machine to Machine (M2M) Communications, Challenges and Opportunities: Final position paper," Technology Strategy Board May 2013.
- [4] Xively - Public Cloud for the Internet of Things (2013). Available: <https://xively.com/>
- [5] Nimbits. (2013). Service for connecting people, sensors and devices on the cloud. Available: <http://nimbits.com/>
- [6] ThinkSpeak. (2013). Internet of Things - ThinkSpeak. Available: <https://www.thingspeak.com/>
- [7] ThingWorx. (2013). The 1st Application Platform for the Connected World. Available: <http://www.thingworx.com/>
- [8] B. Christophe, V. Verdot, and V. Toubiana, "Searching the 'Web of Things'," in Proc. Fifth IEEE International Conference on Semantic Computing, Palo Alto, CA, 2011, pp. 308 - 315.
- [9] T. Luckenbach, P. Gober, S. Arbanowski, A. Kotsopoulos, and K. Kim, "TinyREST – A Protocol for Integrating Sensor Networks Into the Internet," in Proc. Workshop on Real-World Wireless Sensor Networks (REALWSN '05), Stockholm, Sweden, 2005.
- [10] D. Guinard, V. Trifa, T. Pham, and O. Liechti, "Towards physical mashups in the web of things," in Proc. of the 6th IEEE International Conference on Networked Sensing Systems (INSS'09), Piscataway, NJ, USA, 2009, pp. 196-199.
- [11] CONVERGENCE (2013). Available: <http://www.ict-convergence.eu/>
- [12] EBBITS (2013). Available: <http://www.ebbits-project.eu/>
- [13] SOFIA (2013). Available: <http://www.sofia-project.eu/>
- [14] iCore (2013). Available: <http://www.iot-icore.eu>
- [15] R. Goolsby, "Social media as crisis platform: The future of community maps/crisis maps," ACM Trans. Intell. Syst. Technol., vol. 1, no. 1, pp. 1-11, 2010.
- [16] Ushahidi Haiti project Website (2013). Available: <http://www.haiti.ushahidi.com>
- [17] V. Furtado, C. Caminha, L. Ayres, and H. Santos, "Open Government and Citizen Participation in Law Enforcement via Crowd Mapping," Intelligent Systems, IEEE, vol. 27, no. 4, pp. 63-69, 2012.
- [18] DiY Smart Experiences (DiYSE) (2013). Available: <http://dyse.org>

- [19] R. Cyganiak and A. Jentzsch. (2011). The Linking Open Data cloud diagram. Available: <http://lod-cloud.net/>
- [20] R. W. Thomas, D. H. Friend, L. A. Dasilva, and A. B. Mackenzie, "Cognitive networks: adaptation and learning to achieve end-to-end performance objectives," *IEEE Communications Magazine* , vol. 44, no. 12, pp. 51-57, 2006.
- [21] J. O. Kephart and D. M. Chess, "The Vision of Autonomic Computing," *Computer*, vol. 36, no. 1, pp. 41-50, 2003.
- [22] P. Vlacheas, R. Giaffreda, V. Stavroulaki, D. Kelaidonis, A. Somov, V. Foteinos, G. Poullos, A.-R. Biswas, K. Moessner, P. Demestichas, "Enabling Smart Cities Through a Cognitive Management Framework for the Internet of Things", *IEEE Communications Magazine*, vol. 51, no. 6, June 2013.
- [23] Cardone, G., Foschini, L., Borcea, C., Bellavista, P., Corradi, A., Talasila, M., & Curtmola, R. (2013). Fostering ParticipAction in Smart Cities: a Geo-Social CrowdSensing Platform. *IEEE Communications Magazine*, 51(6).
- [24] "Controls drives & automation," Aug. 2012. Available: <http://content.yudu.com/Library/A1ykj7/ControlsDrivesampAut/resources/44.htm>
- [25] Eagle, N., & Pentland, A. (2006). Reality mining: sensing complex social systems. *Personal and ubiquitous computing*, 10(4), 255-268.
- [26] Howe, J. (2006). The rise of crowdsourcing. *Wired magazine*, 14(6), 1-4.
- [27] Lynch, C. (2008). Big data: How do your data grow?. *Nature*, 455(7209), 28-29.
- [28] Mitton, N., Papavassiliou, S., Puliafito, A., & Trivedi, K. S. (2012). Combining Cloud and sensors in a smart city environment. *EURASIP Journal on Wireless Communications and Networking*, 2012(1), 1-10.
- [29] Kramer, J., Magee, J., Self-Managed Systems: An Architectural Challenge in: Briand, L. C., Wolf, A. L., *Future Of Software Engineering (FOSE 2007)*, Minneapolis, MN, USA, pp. 259-268, 2007.
- [30] D. Weyns, B. Schmerl, V. Grassi, S. Malek, R. Mirandola, C. Prehofer, J. Wuttke, J. Andersson, H. Giese, and K. Goschka, On Patterns for Decentralized Control in Self-Adaptive Systems, *Software Engineering for Self-Adaptive Systems, SefSAS II, Lecture Notes in Computer Science*, Springer, 2012.
- [31] Calinescu, R., Ghezzi, C., Kwiatkowska, M., and Mirandola, R., Self-Adaptive Software Needs Quantitative Verification at Runtime. *Communications of the ACM*, 55(9):69-77, 2012.

[32] L-J. Zhang, Q. Zhou. CCOA: Cloud Computing Open Architecture. In the IEEE International Conference on Web Services (ICWS 2009), Los Angeles, CA, July 2009.

[33] R. Chow, P. Golle, M. Jakobsson, et. al., Controlling data in the cloud: outsourcing computation without outsourcing control, in Proceedings of the 2009 ACM Workshop On Cloud Computing security, Chicago, IL, November 2009.

	<p>Javier Poncela received the M.Sc. degree in telecommunication engineering from the Polytechnic University of Madrid, Spain, in 1994 and the Ph.D. degree from the University of Málaga, Spain. He worked in Alcatel Spacio before joining the University of Málaga at the Communication Engineering Department. He has actively collaborated with multinational companies (Nokia, AT4wireless) on formal modeling and system testing in Bluetooth, UMTS and satellite systems. His actual research interests include methodologies for efficient development of complex communications systems, analysis of end-to-end QoS over heterogeneous networks and systems and models for the evaluation of QoE.</p>
	<p>Panagiotis Vlacheas received his Diploma and Ph.D. degree in electrical and computer engineering from the National Technical University of Athens (NTUA). Currently, he is a senior research engineer at the Laboratory of Telecommunication Networks and Services at the University of Piraeus. His research interests are in the areas of radio resource management, reconfigurable and cognitive networks, optimization theory, self-organizing networks, and cognitive IoT.</p>
	<p>Raffaele Giaffreda is the Head of the research group Smart IoT at CREATE-NET. He graduated at Politecnico Torino (Italy) and also holds an MSc in Telecom Engineering from University College of London (UK). He worked in the research departments of Telecom Italia (1994-1995) and British Telecom (1998-2008). He is currently responsible for the Internet of Things research activities in CREATE-NET and he coordinates the EU FP7 Internet of Things project iCore (www.iot-icore.eu).</p>
	<p>Suparna De is a post-doctoral Research Fellow at the Centre for Communication Systems Research at the University of Surrey, UK. Since completing her PhD at the University of Surrey in November 2009, she has been involved in several EU projects such as IoT-A, iCore and IoT.est which investigate various aspects of service provisioning in the Internet of Things domain. Suparna did her BSc. in Physics (Honours) at St. Stephen's College, Delhi University, India and her MSc in Communications Networks and Software at the University of Surrey. Her research interests include discovery and retrieval methods, Web of Things, semantic association analysis and knowledge engineering methods.</p>



Massimo Vecchio received the Laurea degree in Computer Engineering (Magna cum Laude) from the University of Pisa and the Ph.D. degree in Computer Science and Engineering (with Doctor Europaeus mention) from IMT Lucca Institute for Advanced Studies in 2005 and 2009, respectively.

His research background is on computational and artificial intelligence techniques, such as metaheuristics for global optimization and fuzzy logic. During his Ph.D. degree, however, his research interests moved towards power-efficient engineering and application designs for pervasive systems and devices. From October 2008 to March 2010 he worked as simulation engineer at INRIA-Saclay (France). Then, he joined the Signal Processing in Communications group at the University of Vigo (Spain) as a post-doctoral researcher, until September 2012. Recently he moved back to Italy to work as senior researcher at CREATE-NET, mainly in the field of Internet of Things devices and resources virtualization.

Besides, current Massimo's research interests include metaheuristic and stochastic methods for wireless sensor nodes self-localization, node mobility for efficient data collection and power-aware consensus techniques. He is author of one book monograph and co-author of two book chapters, more than 10 journal papers and various conference papers.



Septimiu Nechifor received the Dipl.-Eng. Degree in Electronics Engineering in 1998, the M.Sc. degree in Electronics Engineering in 1999- all from Transilva University in Brasov, Romania. He is currently Research Group Head at Siemens Corporate Technology. He is author or co-author of more than 10 scientific publications in international conferences and journals. He is currently involved on a number of European/National or Industry funded R&D projects (e.g. Internet of Things dynamic configurations, autonomic overlay systems for internet structures). His current research interest is focused on the areas of Autonomic Computing and Networking, Constraints Based Systems, Stream and Big Data analytics.



Raquel Barco holds a M.Sc. and a Ph.D. in Telecommunication Engineering from the University of Malaga. From 1997 to 2000, she worked at Telefonica in Madrid (Spain) and at the European Space Agency (ESA) in Darmstadt (Germany). From 2000 to 2003, she worked part-time for Nokia Networks. In 2000 she joined the University of Malaga, where she is currently Associate Professor. Her research interests include satellite and mobile communication networks, mainly focusing on Self-Organizing Networks.



Mari Carmen Aguayo-Torres received the Ph.D. degree in Telecommunication Engineering from the University of Malaga, Spain, in 2001 (M.S. 1994) with a thesis on Adaptive OFDM. She is currently an Associate Professor in the department of Communications Engineering at University of Malaga. She has published more than 50 papers in international journals and conferences and her main research interests include adaptive modulation and coding for fading channels, generalized MIMO, cooperative communications, OFDM and SC-

	<p>FDMA, crosslayer design, and probabilistic QoS guarantees within mobile cellular, satellite, underwater, M2M and vehicular networks.</p> <p>Dr. Aguayo-Torres is involved in a number of public and private funded projects and actively collaborates with industry, mainly in the field of wireless communications (LTE, LTE-Advanced, WiMax).</p>
	<p>Vera Alexandra Stavroulaki is an assistant professor at the Department of Digital Systems of the University of Piraeus, Greece, and holds a diploma in informatics from Athens University of Economics and Business and a Ph.D. degree in electrical and computer engineering, from NTUA. She has long-term experience in European research and development projects. Her research interests include systems for intelligent environments, cognitive management functionality in various domains, and IoT.</p>
	<p>Klaus Moessner is a Professor in Mobile Communications in the Centre for Communication Systems Research at the University of Surrey, UK. Klaus earned his Dipl-Ing (FH) at The University of Applied Sciences in Offenburg, Germany, an MSc from Brunel University, UK and his PhD from the University of Surrey (UK). His research interests include reconfigurability of different system levels, as well as service provision in the IoT. He is involved in investigation and teaching of mobile service platforms, service oriented architectures, mobile service delivery and service enablers.</p>
	<p>Panagiotis Demestichas is a professor at the Department of Digital Systems at the University of Piraeus, being head of the department since September 2011. He has more than 20 years of experience in national and international R&D projects. His research interests include the design and performance evaluation of wireless and fixed broadband networks, software engineering, service/network management, algorithms, complexity theory, and queuing theory.</p>