Enabling Smart Cities through a Cognitive Management Framework for the Internet of Things

Panagiotis Vlacheas, University of Piraeus
Raffaele Giaffreda, CREATE-NET
Vera Stavroulaki, Dimitris Kelaidonis, Vassilis Foteinos, George Poulios, and Panagiotis Demestichas, University of Piraeus
Andrey Somov, and Abdur Rahim Biswas, CREATE-NET
Klaus Moessner, University of Surrey

ABSTRACT

The Internet of Things (IoT) is expected to substantially support sustainable development of future smart cities. This article identifies the main issues that may prevent IoT from playing this crucial role, such as the heterogeneity among connected objects and the unreliable nature of associated services. To solve these issues, a cognitive management framework for IoT is proposed, in which dynamically changing real-world objects are represented in a virtualized environment, and where *cognition* and *proximity* are used to select the most relevant objects for the purpose of an application in an intelligent and autonomic way. Part of the framework is instantiated in terms of building blocks and demonstrated through a smart city scenario that horizontally spans several application domains. This preliminary proof of concept reveals the high potential that self-reconfigurable IoT can achieve in the context of smart cities.

INTRODUCTION

Urban living already poses major challenges in our daily lives. The United Nations Population Fund forecasts that by 2030 approximately 60 percent of the world population will live in an urban environment, while 27 megacities with greater than 10 million people are anticipated to exist [1]. Therefore, urgent solutions are sought for viable living conditions and sustainable city development. On the other hand, the tremendous development of information and communications technologies (ICT) and the power of the Internet may help us see the light at the end of the tunnel.

Even at present, we can hardly imagine our life without the Internet. With more and more connected objects becoming available, it will be even harder to imagine that the Internet we use today to reach people, find information, and assist us in our daily lives with innovative services will not be extended toward creating value out of such increased availability of connected objects. This article is about helping to achieve that transition from the Internet of People to the Internet of Things (IoT) *for People*, especially in the context of future smart cities.

According to the European Research Cluster on the Internet of Things, IoT is a dynamic global self-configuring network infrastructure where physical and virtual "things" are identifiable and communicate with standard and interoperable protocols. While it is anticipated that IoT will be the great hope of sustainable city life [2], the evidence on how this will be achieved is not so straightforward because of:

- Technological barriers among objects when being used across application domains
- Availability of a wide set of *predefined* services and applications

Within this contextual background we propose a cognitive management framework that aims at showing, besides the widely addressed issue of *how* objects should be connected, *why* and *when* objects need to be connected, and what value they can bring in enhancing existing services and applications.

To this effect, the proposed framework focuses mainly on four axes: how we plan to hide heterogeneity of connected objects, how we ensure resilience (ability to deliver service that can be justifiably trusted in spite of continuous changes) of dynamic service provisioning, how we instruct systems to assess proximity (metric of relevance) between IoT applications and "useful" objects, and how the use of cognitive technologies [3] (ability to dynamically select behavior through self-management, taking into account information and knowledge on the context of operation

0163-6804/13/$25.00 © 2013 IEEE

as well as policies) offers intelligence while minimizing users' intervention. Moreover, some experimental evidence to support the theory behind our proposed framework is provided.

Related Work

The previous section has identified the heterogeneity of objects and the need for resilience as main challenges to be addressed by IoT in the effort for a sustainable smart city, while it has also introduced proximity and cognition as promising factors toward this direction.

In the scope of IoT, the virtualization of objects and the use of semantics are the most important assets for addressing the heterogeneity of objects. The authors in [4] use virtual sensor abstraction in order to hide the complexity in large-scale interconnected sensor networks. The CONVERGENCE project (www.ict-convergence.eu/) introduces a common container for any kind of digital data, including representations of services, people, and real-world objects. The EBBITS project (www.ebbits-project.eu/news.php) aims at transforming every device into a web service with semantic resolution. The CASAGRAS project (www.iot-casagras.org/) suggests the object's identification in the physical layer through the use of ontology instead of representations. However, there is a need for dynamic and intelligent virtual representation of objects, which will allow them to become virtually "always on" in accordance with application requirements and despite changing conditions.

Several efforts have been performed in the field of resilience. In [5], the authors present a taxonomy for threats and countermeasures with respect to dependability and security. The ENISA study in [6] intends to provide end-to-end consideration of resilience through an ontology that encompasses different domains and helps hide the heterogeneity of the underlying infrastructure. In [7], a security framework is presented that comprises mechanisms to achieve intrusion detection and reaction. The work in [8] focuses on enhancing the resilience of supervisory control and data acquisition (SCADA) systems against cyber attacks. However, the majority of these efforts concern separate application domains and rely on human intervention in order to resolve problems. The means to enhance the cross-domain resilience capability of IoT, where a huge aggregation of objects and vulnerabilities exist, in a dynamic and automated way are strongly required.

The study of the similarity between two concepts based on their information presents particular research interest. The effort in [9] presents an estimation of semantic proximity between arbitrary concepts and concept instances, using the taxonomic knowledge concepts. The authors in [10] develop a collaborative and proximity-based mechanism for the clustering of semantic web documents according to their characteristics. Semantic models and a web service for querying are used in [11] to support efficient service discovery, while service selection using proximity-based ontology filtering is proposed. The proximity should be expanded in IoT in order to enhance the description, discovery, and selection of objects based on application requirements, moving from proximity among objects to proximity between users/applications and objects.

The research on cognitive management is mainly focused on resource usage and architectural efficiency in future networks. The effort in [12] introduces cognitive and self-x principles in order to optimize the use of the radio resources and spectrum in future networks. The OneFit project (www.ict-onefit.eu/) comprises cognitive systems and opportunistic networks to achieve higher utilization of resources, lower energy consumption, and reductions in the total cost of ownership in the future Internet. The UniverSelf project (www.univerself-project.eu/) designs a unified management framework for overcoming the management complexity of future networking systems by means of federation and self-management. The introduction of cognitive management in IoT will enable self-management and self-configuration for the virtual representations of objects and their (re)use in the context of dynamic service provisioning.

Summarizing, the above references reveal some issues related to how the current trends address the heterogeneity of objects and the need for resilience, while identifying that proximity and cognitive management should expand toward IoT to help smart city development. These aspects are the motivation for our cognitive framework, introduced in the next section.

Framework Overview

Concepts

This effort proposes a cognitive management framework for IoT that addresses the main issues identified above:

- The considerable number of objects that fall within the scope of any IoT applications and their heterogeneity
- The inherently unreliable nature of such objects and the guarantee of the end-to-end resilience
- The complexity associated with the huge quantities of usable objects in a smart city context

To this end we have identified the need to support end users and application providers with techniques that can automatically select *by relevance* all potential objects discovered to be available for use and able to meet given application requirements. The cornerstone of the framework is the concept of the virtual object (VO). A VO is a virtual representation of any real-world object (RWO), which thus becomes virtually "always on." VOs can be dynamically created and destroyed, and they are dynamic objects since they represent dynamically changing RWOs.

In order to further enhance the resilience capability, as well as offer an automated ability to aggregate VOs to meet application requirements, the framework introduces the composite VO (CVO) concept. A CVO is a cognitive mashup of semantically interoperable VOs (and their offered services) that renders services in accordance with the application requirements.

The cornerstone of the framework is the concept of the virtual object. VOs are virtual representations of any real-world object, which thus becomes virtually "always on." VOs can be dynamically created and destroyed, and they are dynamic objects since they represent dynamically changing RWOs.

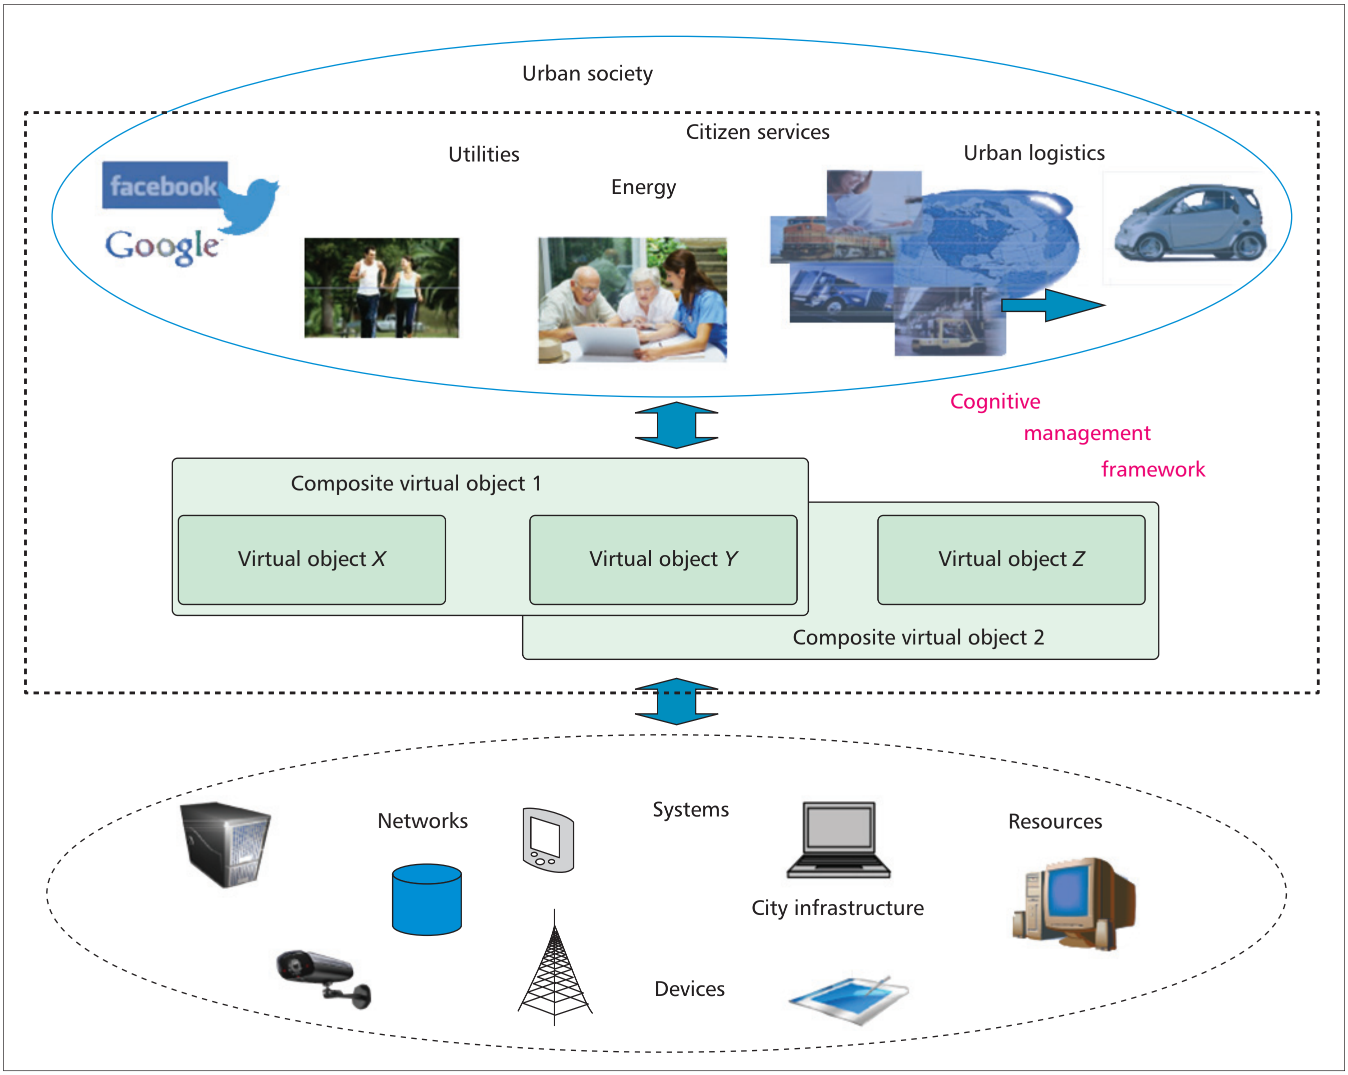

Figure 1. *High-level view of the proposed framework.*

CVOs enable the reuse of existing VOs outside their initial context and domain. Besides this role of fostering aggregation of VOs into more complex services, once created, CVOs are meant to maintain their performance and self-reconfigure as the context changes, giving our framework better levels of resilience/reliability.

As said before, CVOs allow finding the optimal possible way to deliver the application, given both the capabilities offered by VOs and the application requirements. The piece of the jigsaw that completes our framework is the one that provides the functionality of deriving the so-called service logic, which represents the translation of the stakeholder/application requirements into services to be fulfilled by the CVO level. This addresses the need to reduce complexity levels for the users of the platform. Each requested service is accompanied with specific quality of service (QoS) requirements to further address resilience.

This level completes the description of how we equipped our framework to fulfil the three issues identified earlier. The three distinct conceptual levels of our framework, — VO level, CVO level, and service level — are illustrated in Fig. 1.

Cognition and Proximity

The main difference of the proposed framework with respect to existing or past efforts is that cognitive technologies are enabled with the intention of supporting how we address the key issues highlighted above. Cognition is transversal to all three levels. At the VO level, cognition is used for the self-management and self-configuration of VOs in order to maintain a constant link to the relevant RWOs and manage the relevant data flows. More precisely, envisioned cognition includes optimization techniques to select more appropriate links with RWOs, as well as learning techniques to predict problematic links, propose alternate links, and generate knowledge from data.

At the CVO level, cognition is used in order to decide on the VOs and CVOs that should be (re)used to meet the application requirements. CVOs are self-managed self-configurable objects that are created dynamically in an autonomous manner. First, techniques to acquire situation information (e.g., monitoring relevant VOs and their associated RWOs) are crucial, so as to discover VOs/CVOs and exploit their capabilities. Moreover, learning techniques for predicting future situations are also helpful in this direc-

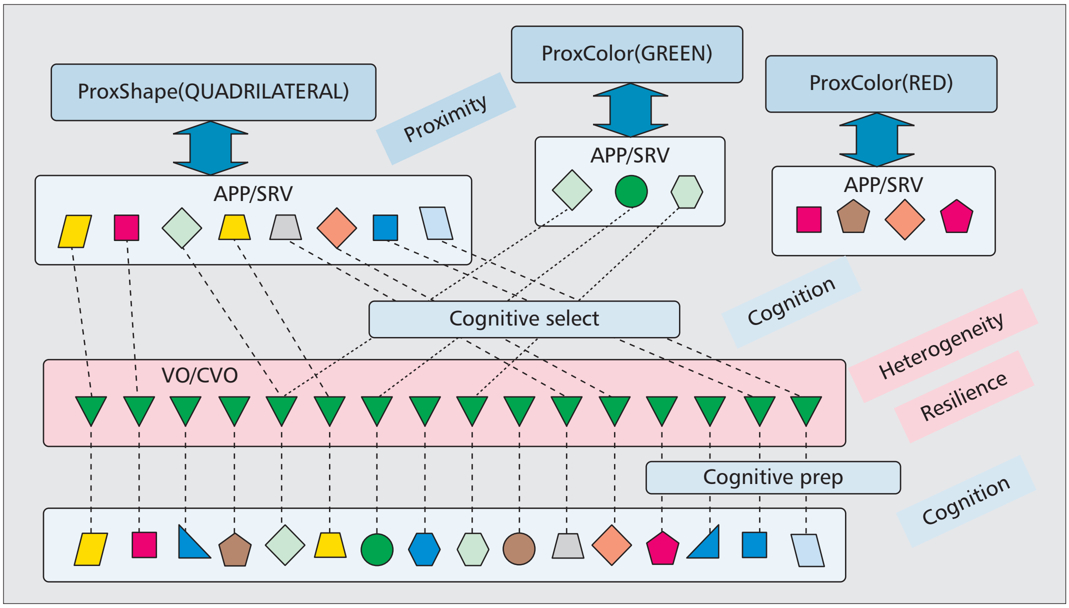

Figure 2. *Technical view of the proposed framework.*

Another vital concept of the framework is the notion of proximity. Proximity denotes the establishment of a level of relatedness between any IoT application and the relevant objects that may be used to deliver this. The actual problem is about assessing how close (hence useful) objects are to users/applications.

tion. Pattern recognition techniques are then used to identify if an already existing CVO can be (re)used for the current application. Otherwise, optimization techniques are exploited afresh for the optimal CVO (re)configuration and selection of CVOs and VOs taking into account user preferences and the current situation. The optimization process is aided by learning techniques for both the prioritization of candidate CVOs or VOs and the exploitation of the knowledge they built in past similar situations. An example of such knowledge derivation is the learning related to the efficiency of CVO (re)configuration.

Cognition at the service level is a prerequisite in order to capture the application requirements and then guide the selection process at the CVO level. Semantic reasoning is used for the automatic translation of a user request to a well formed request for a service, understandable by the system. Learning techniques build knowledge of user preferences and eventually act on behalf of the user. Finally, derivation and management of information related to policies, that is, constraints imposed on the CVO by the administrative domain in which the user is currently located (set of permissible VOs, priorities for the use of certain VOs, security related aspects), are strongly required.

Building cognition should not spoil but needs to coexist with security and privacy features. In the proposed framework, security and privacy are mainly considered through authentication and access control. Authentication provides the means to validate the identity of the user before s/he interacts with the system. Access control is used to regulate access to data and services (through access to the corresponding VOs/CVOs). In this respect, VOs/CVOs are created and managed with their associated policies and access rights, which define when, how, and to what extent the enclosed data/function can be disclosed. The VOs/CVOs can be accessible only by authenticated parties, which are authorized and granted with the appropriate level of access according to their role, even if the VOs/CVOs are distributed across untrusted networks and domains.

Another vital concept of the framework is the notion of proximity. Proximity denotes the establishment of a level of relatedness between any IoT application and the relevant objects that may be used to deliver this. The actual problem is about assessing how close (hence useful) objects are to users/applications. While the current trend is to assess how close objects descriptions or data stored in different databases can be, the differentiation of the framework is toward an effort to bring a much more "application-oriented" approach. Thus, the distance the system wishes to assess is not necessarily only between objects, but between users/applications and objects. Undoubtedly, the proximity is a prerequisite for achieving more automatic and scalable solutions in the cognitive selection of VOs/CVOs.

Figure 2 depicts a technical view of the proposed framework. RWOs are represented by several colors and shapes, highlighting their heterogeneity. One level above, VO/CVO levels are depicted as uniformly colored and shaped objects. Finally, at the top level, several objects are selected to be used by a specific application through the corresponding VOs/CVOs based on proximity. For illustration purposes, the application is considered to aggregate the objects by a criterion related to either their shape or their color. The cognition at the VO level as described above is related to the Cognitive.prep property allowing prepping the data with which to initially enrich the VO. The automatic selection of the most appropriate objects among the available ones for the purpose of an application is represented by the Cognitive.select property. This cognition could be split between the CVO level and the service level. Proximity is used by the Cognitive.select process.

As illustrated by the figure, the concepts of VOs/CVOs are primarily targeted to the abstraction of heterogeneity. This is first achieved by a

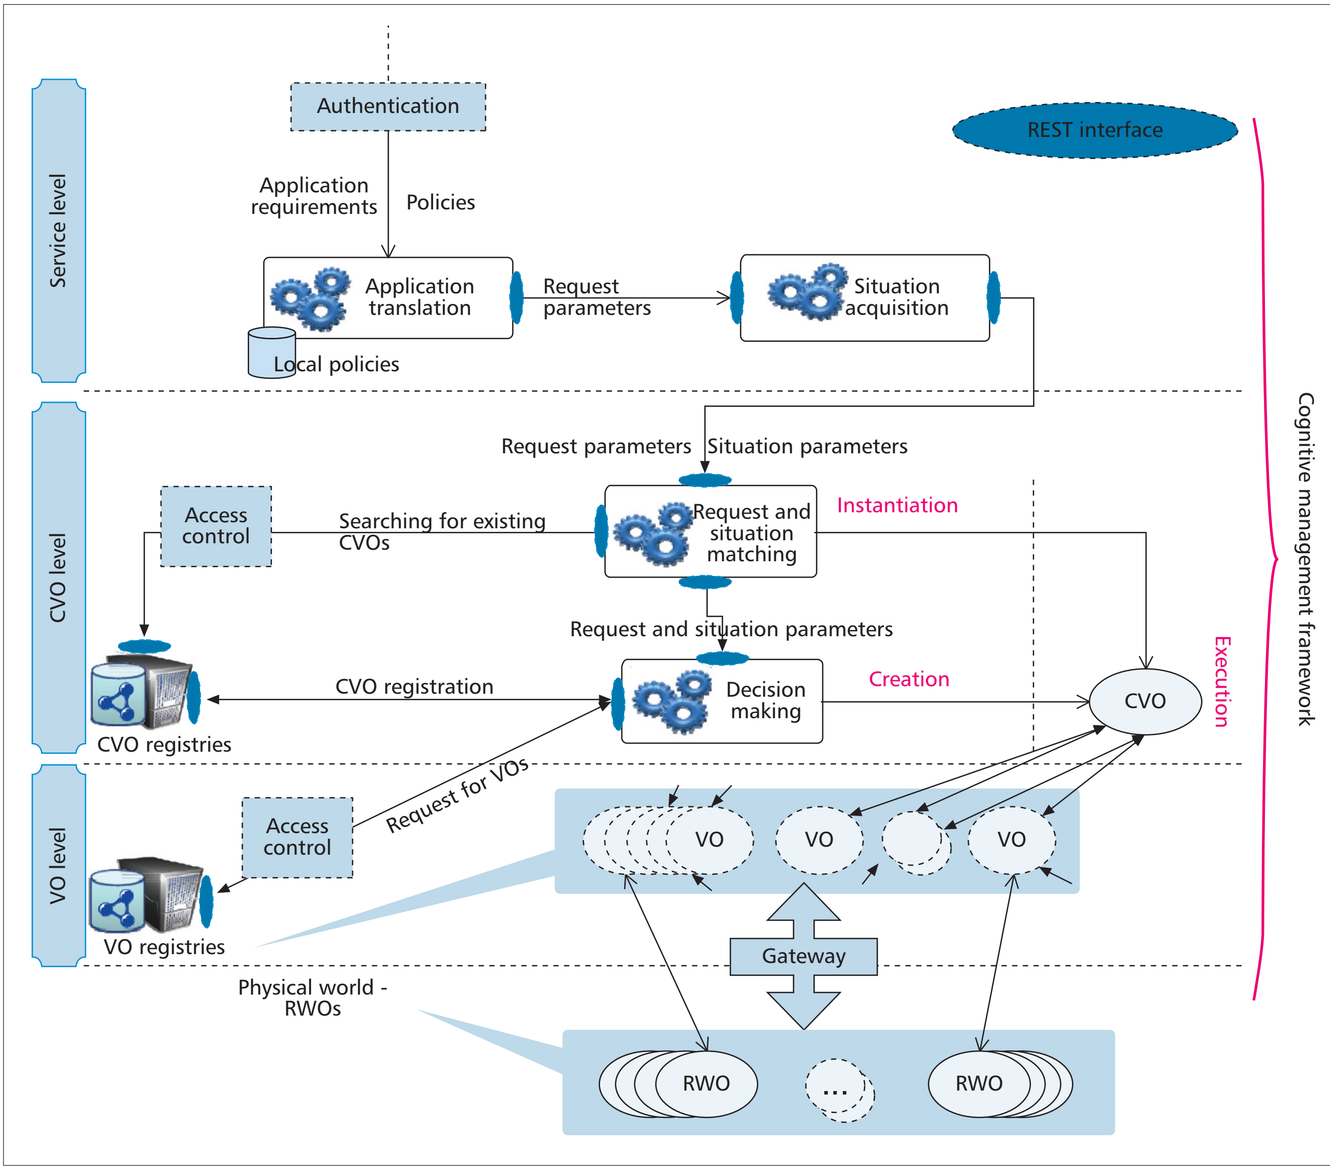

Figure 3. *Framework instantiation by means of building blocks.*

uniform way of representation through semantics. Semantics are used for the specification and high-level description, registration, discovery, and access/invocation of existing VOs and CVOs. As stated before, VOs/CVOs are dynamic, self-aware, self-managed, and self-configurable objects that adapt their own operation, even their own configuration, based on the continuously changing situation related to the underlying objects and overlying application requirements. In such a way, they guarantee the resilience of the application in terms of availability, performance, and dependability, and therefore of the constituent services. Moreover, CVOs abstract the domain dependence by allowing the reuse of VOs outside their initial domain and offer end-to-end resilience.

Framework Instantiation

The previous section illustrated a set of problems we found to be relevant to future smart cities and presented an IoT-based cognitive conceptual framework designed to address them. In this section, we illustrate how such a framework has been instantiated (Fig. 3). For this purpose, a representative set of cognitive functions and components were selected to be implemented as building blocks in the various framework levels, and their role is illustrated hereafter.

Service-Level Components

The users use appropriate interfaces (which are left outside the scope of this specific work) after being authenticated/authorized in order to provide their application requirements and policies, which are then conveyed to Application Translation. This mechanism represents the semantic reasoning for the automatic translation of a user request into a language the system understands. The application requirements are mapped to request parameters (e.g., requested functions), taking into account both the user-provided and local policies.

Then Situation Acquisition is responsible for deriving the situation associated with the request.

The situation is acquired in terms of situation parameters, which enable the system to be aware, as well as to foresee potential changes. The situation parameters may include information such as the time of day, the location of interest, and the available VOs at that time and place. The situation parameters are combined with the request parameters, which helps maximize the reliability of the framework due to more sophisticated and aware decisions at the lower levels of the framework.

CVO-Level Components

Once the request and situation parameters are retrieved, the framework needs to check how it can address those requirements. Request and Situation Matching (RSM) first performs a search in order to discover potentially available, relevant CVOs that can be reused. For this purpose, RSM compares the request and situation parameters with the corresponding recorded parameters in the CVO registry (a registry where logged CVOs are stored) using a satisfaction-rate-based similarity metric (proximity), reflecting the Euclidean distance between them. If there is a matching of a satisfactory percentage, RSM triggers the re-instantiation of existing CVOs in order to be executed. If not matching found, then the CVO should be created from scratch.

In this case, RSM forwards the parameters to Decision Making (DM) with a proposal of available VOs derived based on proximity that might fit the parameters. First, DM interacts with the VO registry (a registry where logged VOs are stored with their semantic annotations, which are used here for an automatic selection process) in order to get information about these VOs (discovery). Then DM selects the appropriate VOs that should be composed for the creation of the CVO (optimization techniques discussed earlier). DM registers the new CVO in the CVO registry and then activates it.

The CVO Registry contains information for each deployed CVO, which is preserved for a specific time period. This information includes:

- CVO identifier
- Request parameters that led to the creation of the CVO
- Situation parameters that represent the context in which the CVO was created
- Access rights of the CVO
- The VOs that comprise the CVO

The CVO registry is implemented and accessible similar to the VO registry, as explained later.

Access Control can impose that only the (C)VOs that can be accessed by a user/application appear after a lookup in the (C)VO registry.

VO-Level Components

The VO is the virtual representation of an ICT object (supporting ICT capabilities such as a smartphone or sensor) that may be associated with one or more non-ICT objects (e.g., persons, places). Both ICT and non-ICT objects are RWOs. The information regarding the VO is semantically enriched by adding contextual information (metadata) through learning techniques and the use of ontology that helps assess proximity in the selection of relevant objects. VOs are implemented as web resources using RESTful Web Services, shortlisted for the wide applicability domain. The VO communicates with the ICT object, either directly or by means of a gateway, using the REST interface over various wireless or wired access technologies. The communication between the ICT object and the gateway in the second case (or even inside a network of ICT objects connected to a VO through a gateway) can be realized by different types of communication technology standards, such as universal plug and play (UPnP), open metering system (OMS), data over cable service interface specifications (DOCSIS), and machine-to-machine (M2M) communication.

Apart from running instances of VOs, the VO level comprises the VO registries, which store the semantically enriched data that are used for the description of the VOs, in order to be available any time from anywhere for the authorized parties. The stored information may include VO identifier, associations with ICT and non-ICT objects, location, offered functions, and access rights of the VO. The semantically enriched data of the VOs are written in Resource Description Framework (RDF) [13], which was designed as a model for representing metadata about web resources, and stored in the RDF graph databases, which implement the VO registries. VO and CVO registries may be distributed across several domains.

Apart from running instances of VOs, the VO level comprises the VO registries, which store the semantically enriched data that are used for the description of the VOs, in order to be available any time from anywhere for the authorized parties.

Case Study and Validation

This section intends to demonstrate the use of the proposed IoT cognitive framework and the cross-application nature of objects in one smart city scenario that horizontally connects several application domains and, more specifically, smart health/smart home/smart living, smart transport, and public safety. Smart cities are actually a real driver for connecting application domains.

Storyline Summary

This scenario refers to how the proposed framework may be exploited to protect and facilitate the daily life in the smart city. Sarah, an elderly woman, has opted for an assisted living service that is provided by a medical center. A doctor, who monitors Sarah's health remotely from the medical center, receives an alarm that Sarah has fainted. An ambulance is instructed to go to assist Sarah. A smart driving application is used by the ambulance to reach Sarah's home as quickly as possible.

Proof of Concept Description

The storyline introduces three different smart city application domains (smart health, public safety, smart transport). The situation in each domain requires a respective CVO:

- The CVO in the medical center for monitoring Sarah's health and the environmental conditions in the smart home
- The CVO in the police department for traffic jam monitoring
- The CVO in the smart vehicle for the smart driving on city streets

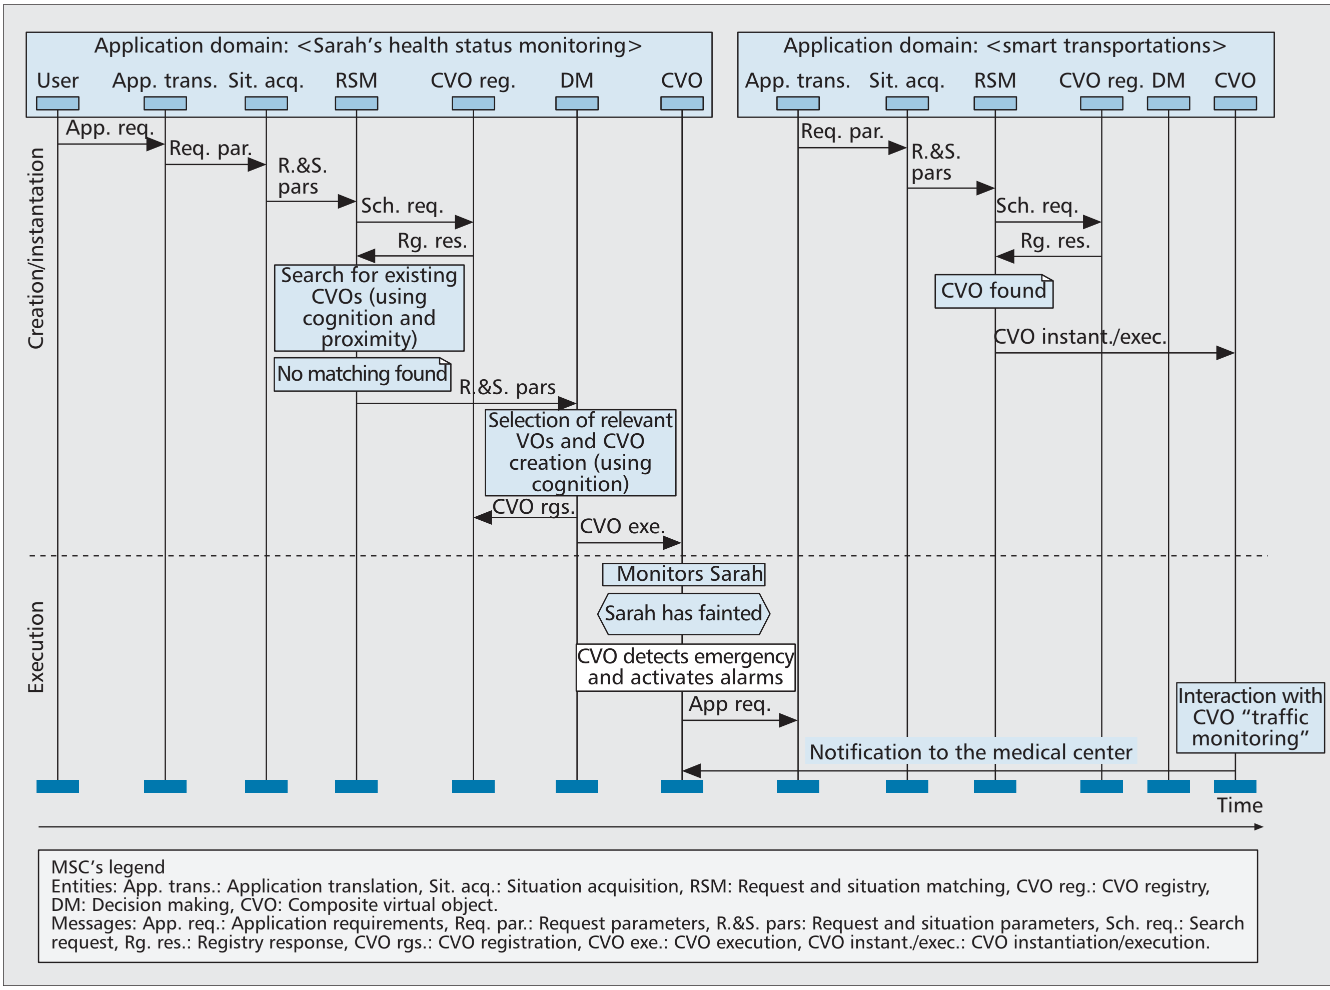

Figure 4. *Use case workflow.*

The use case workflow is depicted in Fig. 4 and is consistent with the description mentioned previously. The requested parameters (e.g., Sarah's pulse, room temperature, alarms [in the medical center and in the smart home] etc.) and the situation parameters (Sarah's location, time, available VOs representing a body pulse sensor, a temperature sensor, a fire alarm, etc.) are extracted by Application Translation and Situation Acquisition to be used for deriving the CVO in the medical center. The scenario assumes that an appropriate existing CVO cannot be found by RSM, so DM selects the appropriate VOs and triggers the creation of a new CVO. The fire alarm is selected to play the role of the alarm in the smart home based on proximity. This selection of an object to deliver an application requirement ("raise an alarm") in an approximate way ("make noise") enhances the resilience of the service offered by the CVO in the absence of more appropriate objects. The CVO is then executed.

When Sarah faints, the CVO becomes aware of the accident, since it receives periodic measurements from the body pulse VO and compares them with a predefined threshold. An alarm is raised by the CVO in both the smart home (fire alarm) and the medical center, the latter being automatically transmitted to an ambulance, potentially with some policies about the seriousness of the incident. The ambulance requests a CVO for smart driving in order to find the best route to Sarah's home and save valuable time. RSM finds and instantiates an existing CVO that matches the request. This CVO receives input data from the Traffic Monitoring CVO in the police department, which combines real-time data from different types of sensors (e.g., cameras, gas sensors, magnetic field sensors [MFSs]) and infers the traffic status in the city's streets. For instance, if the MFSs detect many cars and the carbon dioxide levels are higher than a value, it can infer that there is a traffic jam in the specific area.

The implementation aspects are depicted in Fig. 5. The emulated smart city includes the smart home of Sarah, which is equipped with actuators and sensors hosted by wireless sensor network (WSN) platforms A and B, respectively. The CVO in the medical center uses the Cosm online database service in order to get the sensor measurements from WSN platform B as well as to send the actuator commands to WSN platform A. In the smart city, there are various types of sensors, which are connected through WSN platform C using various access technologies and/or communication protocols (ZigBee, Bluetooth, Wi-Fi, etc.) sharing their data. The CVO in the police department receives data from the distributed platforms of type C in order to infer the traffic status in the city's streets.

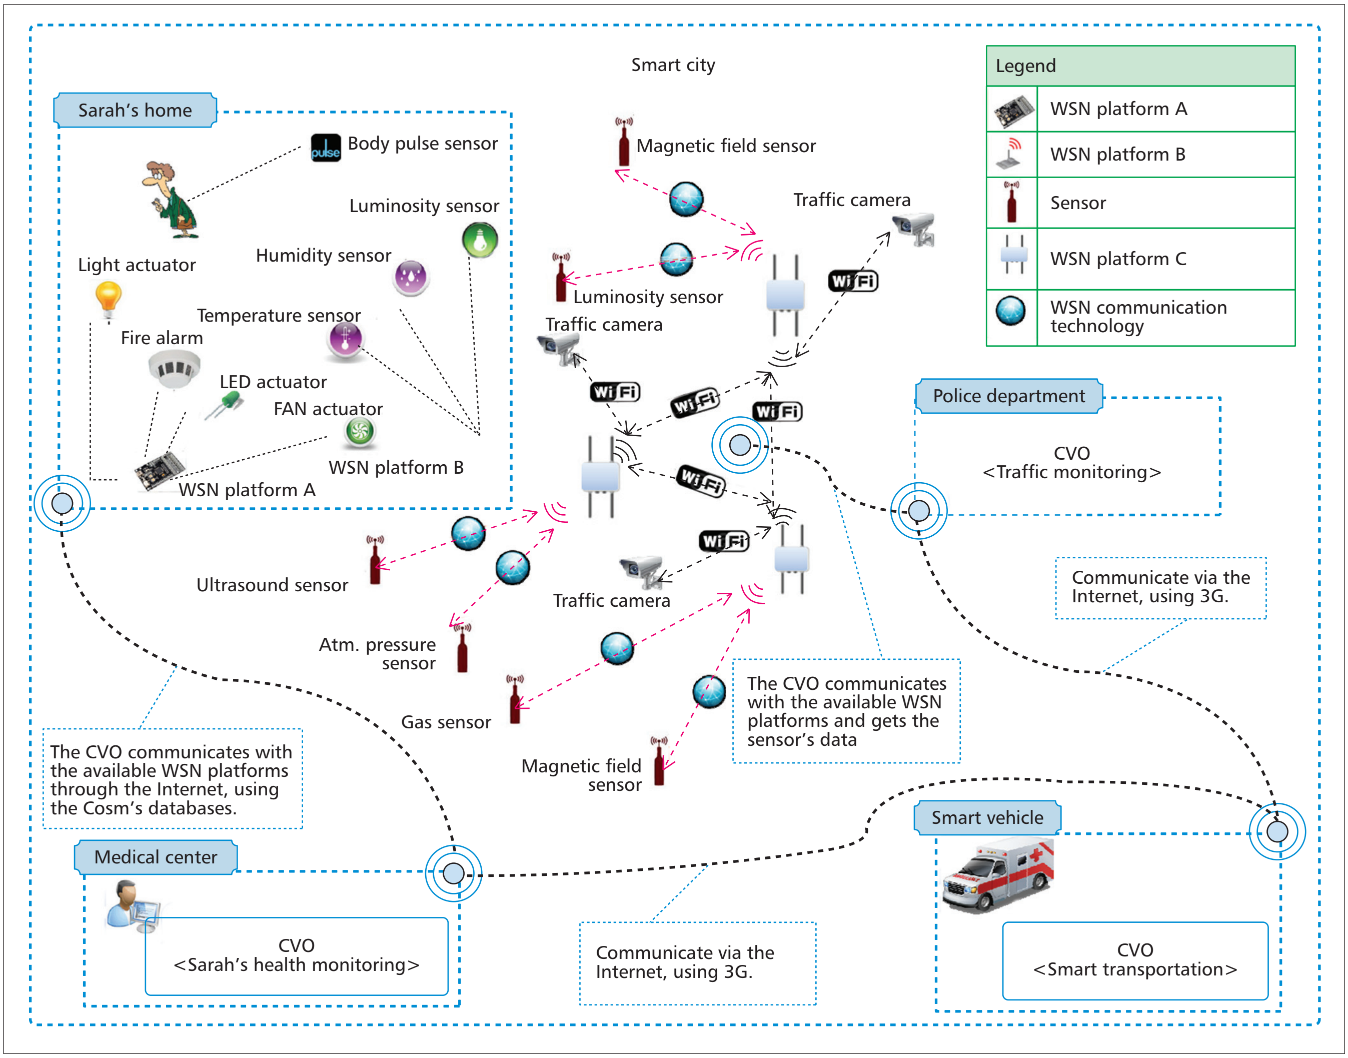

Figure 5. *Implementation aspects of the smart city scenario.*

For the performance evaluation of the framework, we have conducted an indicative experiment. The application requests a service that consists of four functions (which represent the request parameters). The service execution time for different numbers of available objects that are located in the area of interest is presented in Fig. 6 for three cases. The first one corresponds to the lack of the framework with respect to cognition/proximity at the CVO level using an exhaustive search technique for selecting the appropriate objects (case A), the second one to CVO creation through the DM process using the CPLEX algorithm [14] (case B), and the last one to the CVO instantiation (reuse of an existing CVO) through the RSM process based on ranking CVO records according to a satisfaction rate similarity metric (case C). As depicted, the service execution time is decreased through DM, while further time savings are achieved through RSM, showing that cognition/proximity can decrease operational expenditures (OPEX).

The benefits of the framework can be summarized as follows:

- Easy creation and delivery of value-added services from business actors allowing them to increase profits and market share
- Service provision dynamically tailored to the needs of end users, even being able to act on behalf of them
- Decrease capital expenditures for telecom and service providers due to gradual replacement of legacy technology enabled by VO abstraction without disrupting currently delivered services
- Development of innovative cross-domain applications due to CVO abstraction
- Support of large-scale networking through cognition
- Reduced OPEX in terms of time, maintenance, and energy consumption via cognition/proximity
- Decrease in time to market by reusing available VOs
- New business roles as (C)VOs may be owned by different stakeholders

Conclusions

This article presents a cognitive management framework that will empower the Internet of Things to better support sustainable smart city development. Our study revealed the need to address the heterogeneity of objects, while guar-

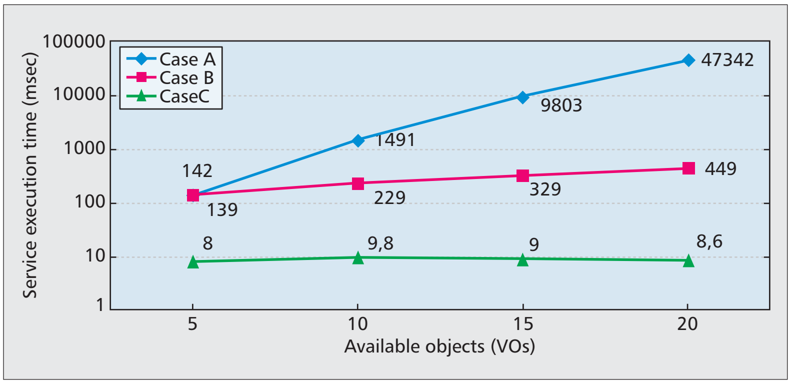

Figure 6. *Service execution time (logarithmic scale): case A: no cognition/proximity; case B: CVO creation through DM; case C: CVO instantiation (reuse) through RSM.*

anteeing the resilience of associated services and their dynamic provisioning. Proximity and cognition were identified as promising solutions that enable applications to use the most relevant connected objects in an intelligent and autonomic way. The framework introduced the virtual object concept as a dynamic virtual representation of objects and proposed the composite VO concept as a means to automatically aggregate VOs in order to meet users' requirements in a resilient way. In addition, it illustrates the envisaged role of service-level functionality needed to achieve the necessary compliance between applications and VOs/CVOs, while hiding complexity from end users. The envisioned cognition at each level and the use of proximity are described in detail, while some of these aspects are instantiated by means of building blocks. A case study, which presents how the framework could be useful in a smart city scenario that horizontally spans several application domains, is also described. Future plans include the implementation of the remaining concepts and a more detailed proof of concept in terms of qualitative/quantitative results.

Acknowledgment

The research leading to these results has been performed within the iCore project (http://www.iot-icore.eu/) and received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 287708.

References

[1] M. Naphade *et al.*, "Smarter Cities and Their Innovation Challenges," *IEEE Computer Society*, vol. 44, no. 6, June 2011, pp. 32–39.
[2] M. Dohler *et al.*, "Smart Cities: An Action Plan," *Barcelona Smart Cities Congress 2011*, Barcelona, Spain, Nov.–Dec. 2011.
[3] J.O. Kephart and D. M. Chess, "The Vision of Autonomic Computing," *IEEE Computer*, vol. 36, no. 1, Jan. 2003, pp. 41–50.
[4] K. Aberer, M. Hauswirth, and A. Salehi, "Infrastructure for Data Processing in Large-Scale Interconnected Sensor Networks," *Proc. Int'l. Conf. Mobile Data Management*, Washington, DC, USA, 2007, pp. 198–205.
[5] A. Aviienis *et al.*, "Basic Concepts and Taxonomy of Dependable and Secure Computing," *IEEE Trans. Dependable and Secure Computing*, vol. 1, no. 1, Jan.–Mar. 2004, pp. 11–33.
[6] ENISA, 2011, Ontology and Taxonomies of Resilience, http://www.enisa.europa.eu/activities/identity-and-trust/technology-for-resilience/ontology/ontology_taxonomies/at_download/fullReport.
[7] M. Chora *et al.*, "Ontology-Based Decision Support for Security Management in Heterogeneous Networks," *LNCS*, vol. 5755/2009, 2009, pp. 920–27.
[8] M. Chora *et al.*, "Ontology Applied in Decision Support System for Critical Infrastructures Protection," *LNCS*, vol. 6096/2010, 2010, pp. 671–80.
[9] C. N. Ziegler, K. Simon, and G. Lausen, "Automatic Computation of Semantic Proximity Using Taxonomic Knowledge," *Proc. 15th ACM Int'l. Conf. Info. and Knowledge Management*, 2006, pp. 465–74.
[10] V. Loia, W. Pedrycz, and S. Senatore, "Semantic Web Content Analysis: A Study in Proximity-Based Collaborative Clustering," *IEEE Trans. Fuzzy Systems*, vol. 15, no. 6, Dec. 2007, pp. 1294–1312.
[11] David Bell, "Proximity Based Semantic Service Discovery for Pervasive Business," *Proc. 2nd Int'l. Wksp. Enabling Service Business Ecosystems*, vol. 530, no. 6, Athens, Greece, 2009.
[12] G. Dimitrakopoulos *et al.*, "Functional Architecture for the Management and Control of Reconfigurable Radio Systems," *IEEE Vehic. Tech. Mag.*, vol. 4, no. 1, Mar. 2009, pp. 42–48.
[13] W3C Rec. 10, Resource Description Framework (RDF): Concepts and Abstract Syntax, http://www.w3.org/TR/rdf-concepts/, 2004.
[14] IBM ILOG CPLEX Optimizer, http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/.

Biographies

PANAGIOTIS VLACHEAS (panvlah@unipi.gr) 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.

RAFFAELE GIAFFREDA (raffaele.giaffreda@create-net.org) is head of the research group eNTIRE at CREATE-NET. He graduated from Politecnico Torino, Italy, and also holds an M.Sc. in telecom engineering from University College of London, United Kingdom. He worked in the research departments of Telecom Italia (1994–1995) and British Telecom (1998–2008). He is currently responsible for research activities on Smart IoT in CREATE-NET and coordinates the EU FP7 IoT project iCore (www.iot-icore.eu).

VERA ALEXANDRA STAVROULAKI (veras@unipi.gr) 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.

DIMITRIS KELAIDONIS (dkelaid@unipi.gr) obtained his Diploma and Master's degree in digital systems from the University of Piraeus in 2008 and 2010, respectively. Currently, he is a research engineer and Ph.D. candidate at the University of Piraeus, Laboratory of Telecommunication Networks and Services, in the area of semantic technologies for future networks. His main interests include the design and development of mechanisms based on semantic technologies for the enhancement of M2M communications in the future Internet era.

ANDREY SOMOV (asomov@create-net.org) is a researcher at CREATE-NET, Italy. He received his Ph.D. degree in electronic engineering from the University of Trento, Italy. His research topic regarded power management in WSNs. In fall 2008 he was a visiting scholar at the University of California, Berkeley, where he conducted research on energy-efficient WSNs. His current research interests include adapting WSNs and cognitive technologies to IoT.

VASILIS FOTEINOS (vfotein@unipi.gr) graduated from the Department of Informatics and Telecommunications of the

National and Kapodistrian University of Athens in 2008 and received his M.Sc. diploma from the Department of Digital Systems of the University of Piraeus in 2011.Currently he is working as a research angineer at UPRC, where he is also pursuing a Ph.D. in traffic engineering for future networks.

GEORGE POULIOS (gpoulios@unipi.gr) received his diploma from the Department of Digital Systems of the University of Piraeus in 2011. He is currently an M.Sc. student in techno-economic management and security of digital systems in the Digital Systems Department of the University of Piraeus. Currently, he is a researcher at UPRC, involved in several integrated projects. His research interests are in the area of machine learning and knowledge engineering for network management and service composition.

ABDUR RAHIM BISWAS (abdur.rahim@create-net.org) leads the IoT research group in CREATE-NET. He has worked on many European projects as project manager (iCore), technical group leader (EUWB), and EC cluster leader. He has acquired experience throughout studying/working in Bangladesh (B.Sc.), Sweden (M.Sc.), Germany (Dr.-Ing), Spain (visiting researcher), and now Italy. He is involved in activities related to IoT and cognitive technologies, as steering committee member (CROWNCOM), leader of IERC AC8, co-chair of SelfIoT, and a member of the management committee of COST (TERRA, IC0902).

KLAUS MOESSNER (K.Moessner@surrey.ac.uk) is a professor in the Centre for Communication Systems Research at the University of Surrey, United Kingdom. He holds a Dipl-Ing degree (Offenburg, Germany), an M.Sc. (Brunel, United Kingdom), and a Ph.D. (Surrey). His research interests include reconfigurability of different system levels, service platforms, cognitive radio, IoT, and smart city deployments. He has led the smart transport activities in the OUTSMART PPP project.

PANAGIOTIS DEMESTICHAS (pdemest@unipi.gr) is with the University of Piraeus, Department of Digital Systems, being head of the department since September 2011. He has been actively involved in a number of national and international R&D programs. Currently, he serves as the deputy leader of the Unified Management Framework work package in the ICT/FP7 UniverSelf Project, and he participates to the IoT/FP7 iCore. He also served as Project Coordinator of the ICT/FP7 OneFIT Project (07.2010-12.2012).