Scalable Computing: Practice and Experience
Volume 13, Number 2, pp. 139–147. http://www.scpe.org
ISSN 1895-1767

A COGNITIVE MANAGEMENT FRAMEWORK TO SUPPORT EXPLOITATION OF THE FUTURE INTERNET OF THINGS

GIANMARCO BALDINI[*], RANGA RAO VENKATESHA PRASAD[†], ABDUR RAHIM BISWAS[‡], KLAUS MOESSNER[§], MATTI ETELAPERA[¶], JUHA-PEKKA SOININEN[‖], DSEPTIMIU-COSMIN NECHIFOR[**], VERA STAVROULAKI [††] AND PANAGIOTIS VLACHEAS[‡‡]

Abstract. In this article, a cognitive management framework is proposed for ensuring exploitation of the Future Internet of Things (FIoT). Cognitive systems offer self-x and learning. A cognitive system has the ability to dynamically select its behavior through self-management/awareness functionality, taking into account information and knowledge on the context of the operations as well as policies and including the generation of the context itself. The framework is based on the principle that any real world object and any digital object that is available, accessible, observable or controllable can have a virtual representation in the Future Internet, which is called Virtual Object (VO). Basic VOs can be composed in a more sophisticated way by forming Composite VOs (CVOs), which provide services to high-level applications and end-users. The described paradigm is applied to various applications scenarios: smart home, smart office, smart city and smart business. This paper presents some background in IoT, identifies the requirements and challenges, and sets the directions that should be followed.

Key words: Cognitive Systems, Internet of Things, Virtual Objects, Wireless Communications

1. Introduction. The "7 trillion devices for 7 billion people" paradigm [1], yields that the handling of the amount of objects that will be part of the Internet of Things (IoT) requires suitable architecture and technological foundations. The Internet-connected sensors, actuators and other types of smart devices and objects need a suitable communication infrastructure. While several research projects (e.g., IoT-A [2], CASAGRAS2 [3]) have set out to define architectures or reference models to ensure interactions and facilitate information exchange, there is a significant lack in terms of management functionality and means to overcome the technological heterogeneity and complexity of the pervasive networks in terms of exploitation effectiveness. This is essential for the IoT, in order to enhance context-awareness (by being able to match continuously an evolutionary demands of client applications against an unreliable connection and representation quality of real world objects), provide high reliability (through the ability to use heterogeneous objects in a complementary manner for reliable service provision), energy-efficiency (through the selection of the most efficient and suitable objects from the set of heterogeneous ones, and, in general, through the optimal management of a large population of resource constrained devices) and security in these distributed networks of cooperating objects. The sheer numbers of objects and devices that have to be handled and the variety of networking and communication technologies, as well as administrative boundaries that have to be supported do require a different management approach. The idea is to enable seamless and interoperable connectivity amongst heterogeneous number of devices and systems, hide their complexity to the user while providing sophisticated services and applications [4].

In response to the requirement of overcoming technological heterogeneity this paper proposes a cognitive management framework. The proposed framework aims to provide the means to realize the principle that any real world object and any digital object, which is available, accessible, observable or controllable, can have a virtual representation in the IoT. This means that the functionality or features offered by any kind of object can become part of composite functionality/features, which will be reusable in the context of sophisticated application/service provision in the IoT.

Moreover, the aim of the framework is to provide the foundations, architecture and functionality for a cognitive support paradigm for the IoT, for continuous sensing of client applications, own environment and real world variations. A cognitive system has the ability to dynamically select its behavior (managed systems configuration), through self-management/awareness functionality, taking into account information and knowledge (obtained through machine learning) on the context of operation (e.g., internal status and status of

[*]JRC, Belgium (gianmarco.baldini@jrc.ec.europa.eu).
[†]TU Delft, The Nederlands (R.R.VenkateshaPrasad@tudelft.nl).
[‡]Create Net, Italy (abdur.rahim@create-net.org).
[§]University of Surrey, UK (K.Moessner@surrey.ac.uk).
[¶]VTT, Finland (Matti.Etelapera@vtt.fi).
[‖]VTT, Finland (Juha-Pekka.Soininen@vtt.fi).
[**]Siemens CT, Romania (septimiu.nechifor@siemens.com).
[††]UPRC, Greece (veras@unipi.gr).
[‡‡]UPRC, Greece (panvlah@unipi.gr).

environment), as well as policies (designating objectives, constraints, rules, etc.).

In the light of the above, cognitive technologies constitute an efficient approach for addressing the technological heterogeneity and obtaining context awareness, reliability and energy efficiency. Cognitive technologies have been applied to the management of diverse heterogeneous technologies (e.g., wireless access, backhaul/core segments). The proposed framework applies this successful paradigm for solving problems that are particular to the IoT. Therefore, new IoT-oriented cognitive functionality will be provided, which will be part of the service layer of the Future Internet. A cognitive system consists of the cognitive engine (offering intelligence and service capabilities) and the reconfigurable/managed part, which is technology specific. The engine interacts with the managed part and with other engines. Each managed part is directly controllable by one engine (i.e., other engines have to interact with the managing cognitive engine in order to affect a specific managed resource). Through this approach there is the accomplishment of the abstraction of the technological heterogeneity, which leads to the removal of the sector specific boundaries. Of course, the realization of above described cognitive capabilities relies intensively on the support of autonomic capabilities on both thing level and support platform level. Context awareness is inherent in the model, while policies and decision-making (part of autonomic features) can be oriented to address the targets of enhanced reliability and energy-efficiency.

Additionally, the proposed framework addresses security, resilience and user privacy issues, which are vital for the Future Internet, though a policy management approach where access to data and resources is regulated by policies and access levels (also common with current practices in autonomic computing and networking).

From the users/applications perspective, three concepts - IoT, ubiquitous computing, and ambient intelligence - aim at delivering smart services (where smartness reflects the accuracy of context sensing, service matching, platform use and time wise capability) to users [5]. A part of the smartness relies on context awareness, e.g., service provision according to the needs that exist at the place, time and overall situation. This is not all. At a societal level, smartness also requires that the needs of diverse users and stakeholders are taken into account both design time (templates) and run-time (learning capability). Stakeholders can be the owners of the objects and of the communication means. Different stakeholders that are part of the current Internet milieu and they will be part of Future Internet, have interests that may be adverse to each other and their very own criteria on how objects could be used and should be accessed. Clark et al. [6], calls this process the tussle and any Future Internet framework should be able to accommodate such tussle to support a smooth evolution of Future Internet of Things, as expected to be a world of competing and conflicting contexts and demands. So a key challenge that needs to be tackled includes the handling of the diversity of information while respecting the business integrity, the needs and rights of users and of the various stakeholders.

The approach presented in this paper aims to overcome the issues above by bringing further intelligence in the Internet of Things. The remaining part of the paper is organized as follows: Section II describes the proposed approach to address such challenges through a cognitive management framework based on the concept of virtual object. Section III describes the Security and Privacy aspects. Section IV describes the application of the cognitive management framework to two scenarios: smart home and smart office. Section V concludes the paper and provides future directions in this research area.

2. Cognitive Management Framework for Future Internet of Things.

2.1. The framework. The proposed framework is targeted to concealing technological heterogeneity and for satisfying the requirements of different users/stakeholders so as to meet the objectives for context awareness, reliability, and energy efficiency. Additionally, security will be a primary concern and an important property at all levels of the cognitive framework. The framework comprises three main levels of enablers, which are reusable to various-diverse applications.

In each level there are scalable fabrics, which offer mechanisms for the registration, look-up and discovery of entities, and the composition of services.

Cognitive entities at all levels provide the means for self-management (configuration, healing, optimization, protection) and learning. In this respect, they are capable of perceiving and reasoning on their context (e.g., based on event filtering, pattern recognition, machine learning), and of conducting associated knowledge-based decision-making (through associated optimization algorithms and machine learning).

Through such features the proposed framework constitutes an open networked architecture encompassing highly intelligent (i.e., adaptive, knowledge based, eventually proactive, etc.) software.

The virtual objects (VOs) are primarily targeted to the abstraction of technological heterogeneity. VOs accomplish their role through the cognitive management and handling of real-world or digital objects (e.g.,

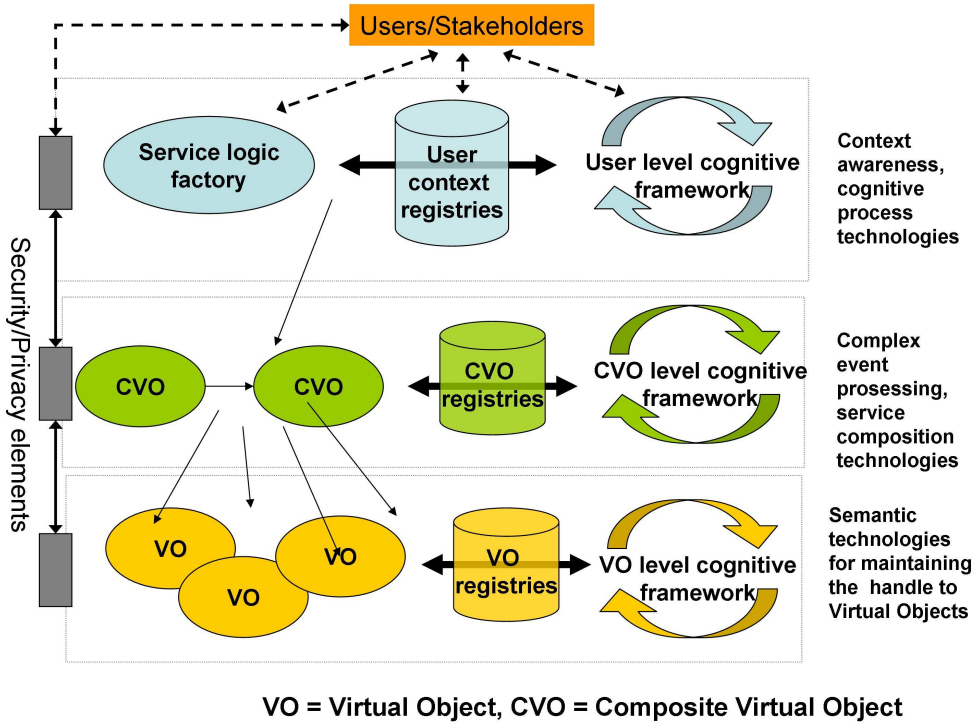

Fig. 2.1: Layers of the Cognitive Management framework

sensors, actuators, devices, etc.). VOs are cognitive virtual representations of real-world objects and/or digital objects.

User/stakeholder related objects will convey the respective requirements. The entities will be capable of detecting human intentions and behavior, inferring, and eventually acting on behalf of the users. In this respect, there is seamless support to users, which is in full alignment with their requirements (the learning capabilities of the cognitive entities of this layer will be applied for acquiring knowledge on user/stakeholder preferences, etc.). Capabilities for governing the entities will also be included (through any type of interaction - multi-modal interactions)

Composite virtual objects (CVOs) will be using the services of virtual objects. A CVO is a cognitive mash-up of semantically interoperable VOs that renders services in accordance with the user/stakeholder perspectives and the application requirements.

The concept of VOs is not new. Object-oriented (OO) approaches have been used in computer programming for decades and distributed objects are used in Object-oriented middleware applications in the Web 7. The intention is not to create new digital representations/objects, but to combine previous concepts with cognitive management mechanisms to create and maintain dynamic, intelligent virtual representation of real world/digital objects that can enhance the Future Internet.

As already introduced, the framework comprises three layers of cognitive components, which are depicted in Figure 2.1.

The first cognitive management layer (*VO level cognitive framework*) is responsible for managing the VOs and for the abstraction of the technological heterogeneity. Real-world or digital objects (e.g., sensors, actuators, devices, etc.) are represented in the first layer through VOs. In current practice and standardization the most prevalent solution is considered the use of RESTful services, based on experienced capabilities of things. The management layer is responsible for the VO lifecycle (i.e., creation, update, destruction) and to address the heterogeneity by defining the logical links among VOs. For example the container transported by a truck is a VO as the truck itself. A tracking device on the truck (with GPS and communication terminal) is also a VO.

The second cognitive management layer (*CVO level cognitive framework*) is responsible for composing the VOs in Composite VO (CVO). CVOs will be using the services of VO to compose more sophisticated objects.

A CVO is a cognitive mash-up of semantically interoperable VOs that renders services in accordance with the user/stakeholder perspectives and the application requirements. For example, the combination of the trucks, the transported goods and the tracking device is represented in the cognitive framework as a CVO.

The third level (*User level cognitive framework*) is responsible for interaction with User/stakeholders. The cognitive management frameworks will record the users needs and requirements (e.g., human intentions) by collecting and analyzing the user profiles, stakeholders contracts (e.g., Service Level Agreements) and eventually acting on behalf of the users. In this respect, there is seamless support to users, which is in full alignment with their requirements (the learning capabilities of the cognitive entities of this layer will be applied for acquiring knowledge on user/stakeholder preferences, etc.). Capabilities for governing the entities will also be included (through any type of interaction - multi-modal interactions).

An alternative representation of the framework is described in Figure 2.2, where the lowest level is connectivity level and it is composed of real world objects. These objects may or may not be communication enabled. Then we define the VOs and CVOs as above. These virtual objects have cognitive functionalities depending on their capability sets. The right composition of abstracted virtual objects is done at this level (i.e., VO level). The topmost level is the service level (may also be called as User level since users could also influence the way services are configured). Based on the required service the CVO are composed on the fly. The user requirements, the context and the analysis of the available resources contribute to the dynamic service composition and orchestration to offer the right services at the right time to the user.

The next section describes more in detail the lifecycle of the VO/CVOs and the dynamic composition of services and applications.

2.2. Virtual object lifecycle and dynamic composition of services. Since they have to represent dynamically changing real world objects, VOs and CVOs should be dynamically created, updated and destructed. More importantly, the services provided by VO and CVO have to be dynamically composed to support the users needs in function of space, time or context [9]. For example: provisioning of communication services and data may be different for Healthcare services during an emergency crisis (i.e., the aftermath of an earthquake) or during routine operations.

The proposed framework encapsulates the support for maintaining and exploitation of VOs and support for transferring the incentives of stakeholders to CVOs. We can say that the framework (and support implementation platform) has to support four levels of cognition. First, the support platform has to keep the continuous link to the real world. Secondly, it has to react to the service context changes of the VO and CVO. For example, the consequences of an earthquake are that VOs representing base stations or routers may be destroyed or provide degraded services (i.e., lower data rates). Thirdly, it must be able to identify, understand, and learn from changes in the context. For example, in normal situations public safety officers may not be allowed to access sensitive information on civilians (i.e., medical conditions) but in emergency crisis, framework must provide the access to these data. Fourthly, it must be able to resolve conflicts between different infrastructures (i.e., set of CVOs) or the tussle described in the introduction. For example, enterprises are always looking for information or resources to undermine the competition from other enterprises. Users would like to use connectivity resources without paying for them (e.g., Skype), while telecom providers would like to maximize the revenues.

The cognitive management framework will also control the lifecycle of the VOs and CVOs: their creation, destruction, and update. While (real world or digital) objects may be owned (controlled) by a particular stakeholder, the VO can be owned (controlled) by particular service providers. And in turn, CVO may be owned (controlled) by yet another provider who adds value by combining different virtual objects and providing these combinations to users. This hierarchical structure leads to a rather complex eco-system, but it opens new opportunities for all stakeholders. Furthermore, the cognitive management system will ensure that the complexity of this eco-system will be well hidden from the different players and stakeholders. The proposed life-cycle and the composition of services and applications are depicted in Figure 2.3.

With reference to Figure 2.3, the following steps describe the lifecycle:

1. A new entity is discovered in the area managed by a specific instance of implementation platform. For example: a PC is switched-on, authenticated and connected to the network. The PC is registered in the VO registry of instance with the related features.
2. In a similar way, CVOs are created by analysis of the VO registry or directly created by composite objects in the real domain (e.g., a taxi).
3. At the user level, an instance authenticates a new user and his set of preferences like type of information he is interested in, type of used terminal and so on. The instance periodically records the users context

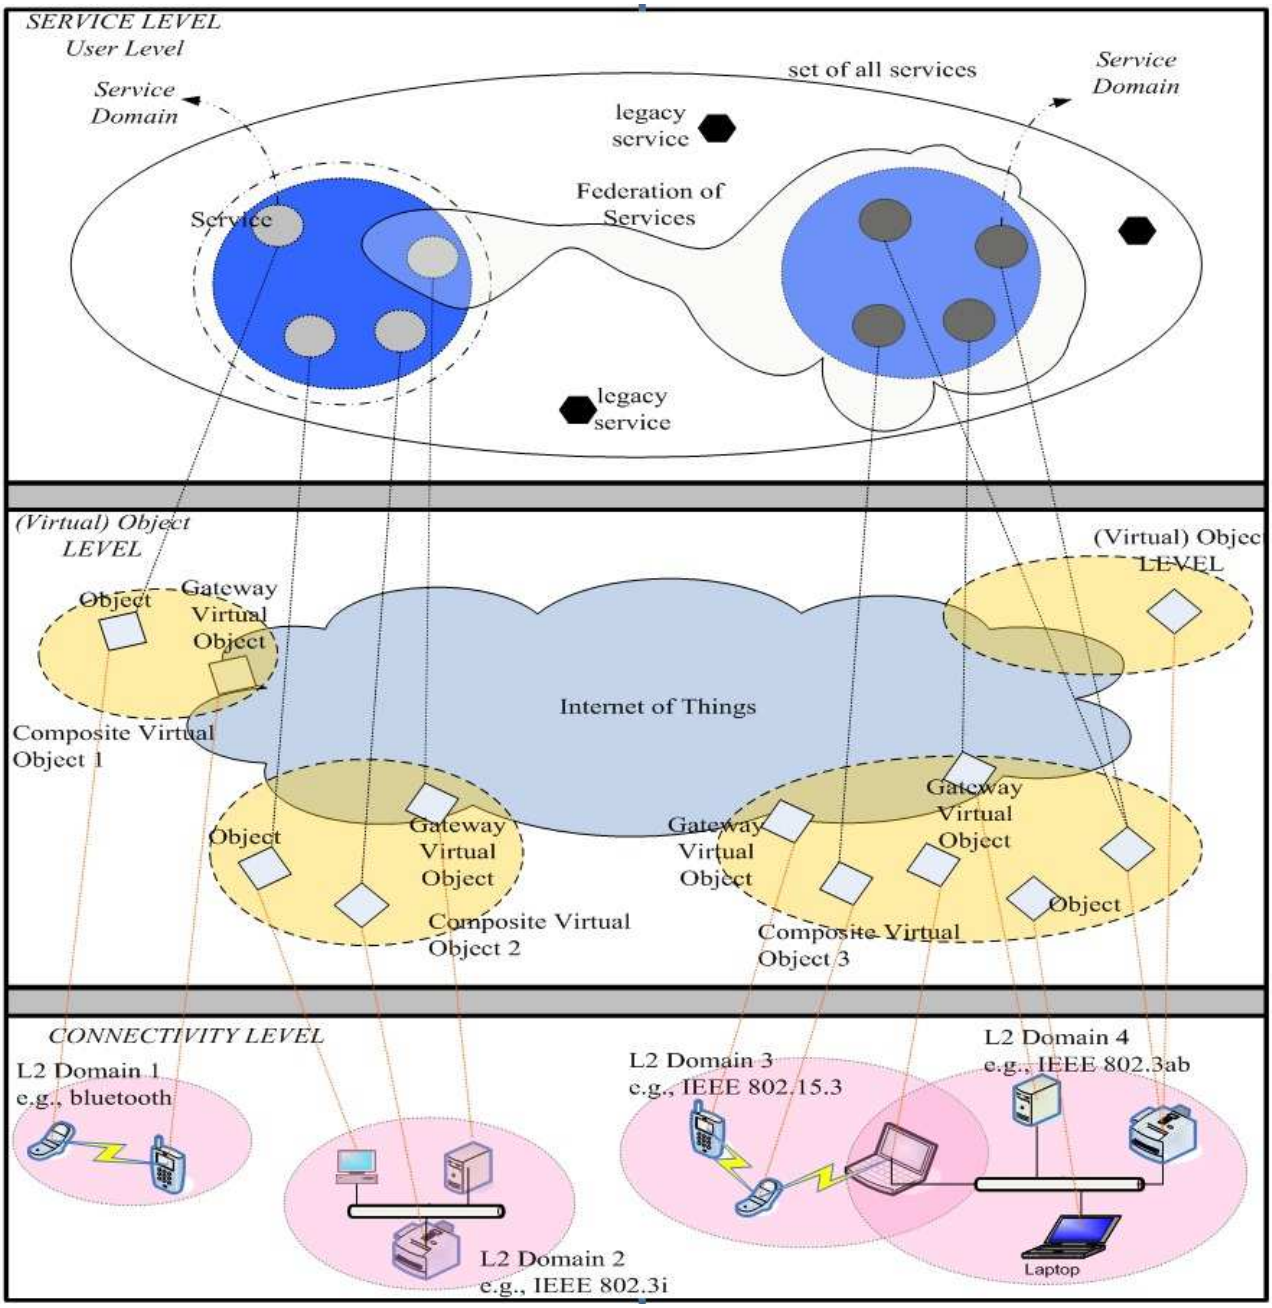

Fig. 2.2: Connect Physical objects, VO and Services

and location.

4. At the service level, the instance correlates the services with the CVO/VO present in the registry. For example, the taxis with credit card payment systems (CVO) present in a specific urban areas, which are available to provide transportation services.
5. Any application can use an instance to access services, information on CVO/VO and users data (see section III for security aspects). For example: a taxi booking application can record the users needs in an area (e.g., need for transportation with credit card) and their location. Then, the application matches the needs with the available CVO (e.g., taxi with credit card services). Specific needs can also be addressed through supported semantic modeling, reasoning or machine learning capabilities. For example: a taxi with a large trunk for suitcases.

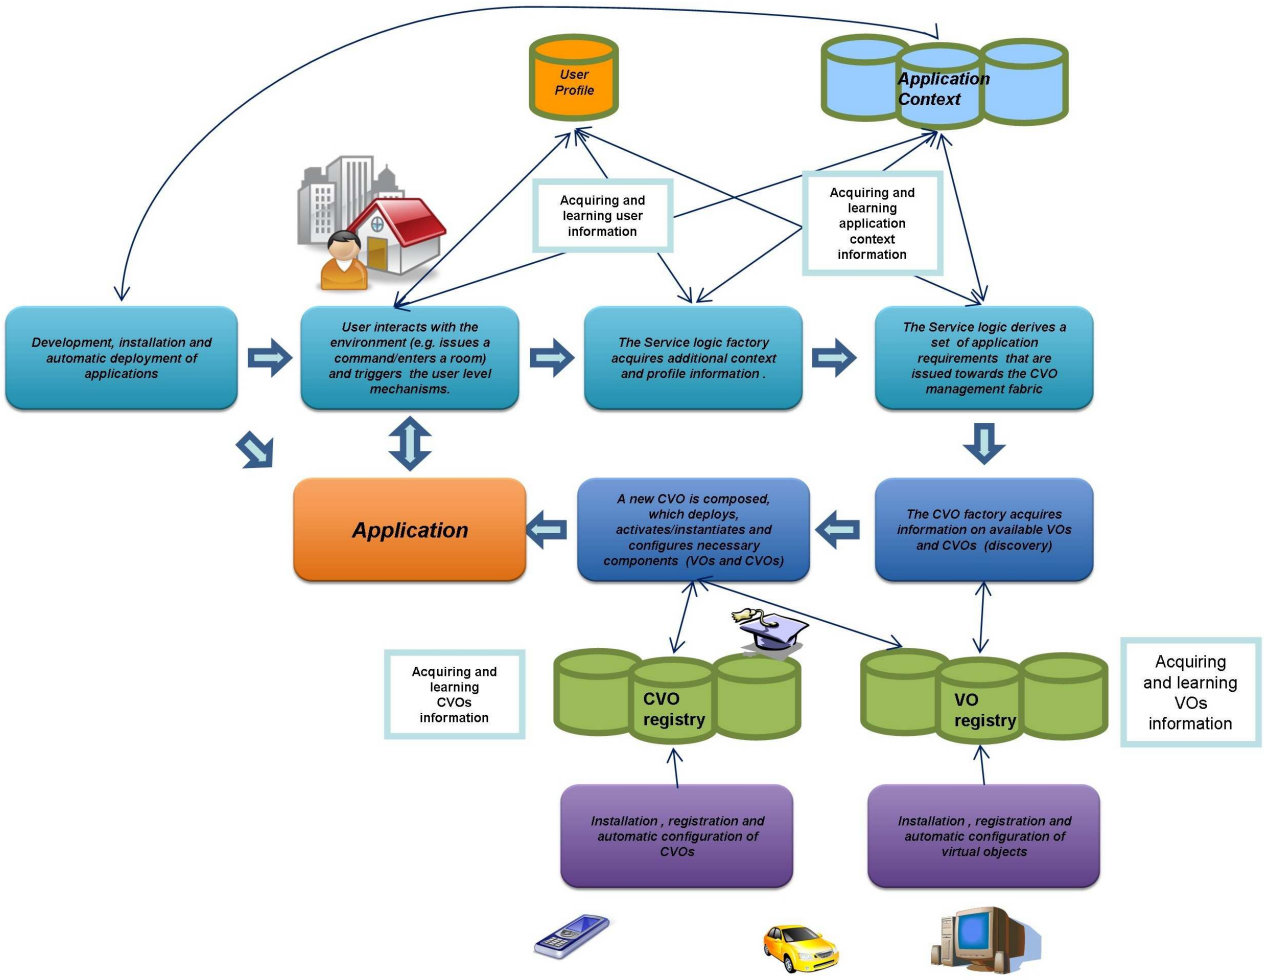

Fig. 2.3: A platform instance life-cycle and services composition

6. An instance can also be used to support the deployment and activation of the application on the users terminals, because records the capabilities of terminals in the CVO/VO registries. It can support a more efficient deployment and tailoring for specific application, beyond the security requirements described in the next section.

3. Security and Privacy Aspects. As described in the previous sections, the proposed framework has to represent (through the VO concept) any type of real object, which may or may not already be connectable to the internet. Real objects could provide privacy sensitive data. In addition, the combination of data from multiple real objects into a new VO could be privacy sensitive. A real object could also provide access to control a particular system (i.e., actuator) for which some kind of access control is required. So the framework has to address the security and privacy issues to enable the dynamic creation of VOs and re-use of real objects and VOs for providing reconfigurable services. The challenge with respect to security and privacy is to integrate novel privacy and security techniques right from the start such that security and privacy would not become an afterthought or add-on feature. Another challenge is to integrate existing legacy systems, which may have proprietary security systems.

The framework is based on the concept that access to data, resources and services represented by VO/CVO must be regulated through a sticky policy management approach [8]. The underlying notion behind Sticky Policy is that the policy applicable to a piece of data travels with it and is enforceable at every point it is used. Users will therefore be able to declare privacy statements defining when, how and to what extent their personal information can be disclosed.

Each VO should contain the following information: 1) the available resources, 2) the access level rights and 3) associated policies valid for this object. This information will be added to the VO in the creation phase, which also includes authentication. For example: a GSM terminal can become a VO, when it is authenticated by the GSM network. In most cases, real-object will not have an authentication mechanism and a security agent must be defined to implement the authentication. In other cases, specific proprietary security systems are already present and a security mediation layer must be defined.

A policy can have multiple rules, which define what resources can be accessed and managed (i.e., used) by a specific access level right. Policies can be defined independent by the existence of associated virtual objects. Note that there are different levels and types of access rights: create, read, modify. Under specific rules, access rights can be delegated: a virtual object can acquire for a specific time, or space or context the access rights of another virtual object. This is particularly useful for the creation of automatic agents. The access level rights can also be used in the ontology and lookup mechanism in the CVO/VO registries: if a virtual object has higher access rights than the entity accessing the dictionary, the virtual object will not appear. Another interesting approach that needs to be considered is claim based access control [10]. In claims based access control access to a resource is provided based on one or more claims. Examples of claims are the proof that "I am over 18 years old", "I am an employee of this company", "I am a guest at this hotel". The enforcement of policies will guarantee privacy of data because data cannot be accessed by entities with the insufficient access right. Data can be protected by additional cryptography services: an entity may have the access rights to access a specific data, but this may be encrypted. Then the entity may also need access to a cryptographic service as well. Policies will also support the concept of Trust Management and reputation scheme. Access rights can be removed or increased by a central authority. For example, an entity may have decreased access rights if there was a security breach. Or, an entity may have increased access rights in occasion of a specific context: for example, public safety officers may have access to sensitive information (e.g., building plans) in case of a crisis management.

Figure 3.1 provides a pictorial view of the security framework and functions. When a VO is authenticated and stored in the VO registry, its access levels and related policies are defined. Users data is also stored as a VO with specific access levels to preserve the privacy of the user. The existing access policies are stored in a distributed policies database, which also records relevant global and local regulations. For example: radio frequency spectrum regulations to regulated use of spectrum in cognitive radio networks or privacy regulation to regulate access to data. The security and policy management functions regulate the access to CVO/VO on the basis of the user access levels and user context. In most cases, existing security frameworks are already in place: for example the authentication functions of GSM/UMTS networks. In these cases, security gateways are required to mediate the necessary information. The main challenge for the implementation of these concepts in the proposed architecture is scalability. A security framework like PKI can handle a specific number of objects/entities but the ICT call requires the management of thousands of objects. Functions like audit and accountability are quite resources consuming if applied to very large networks. A hierarchical approach could be proposed, with hierarchical domains based on geography or contexts. A specific organization can decide to provide only an interface with a subset of VOs to the rest of the Future network.

4. Application Scenarios. This section describes some a set of application scenarios in order to highlight the real life value of the envisioned technology. Cognition and CVOs are cross-domain, and aim to overcome current interfacing efforts. Ranging from home and office domotics use up to industrial infrastructure, the cross-domain aspect generates high volumes of specific developments, often erroneous or inflexible. In computer industry the plug-and-play capabilities have been received as a relevant improvement in terms of functionality and natural use. This fact was based on standardized interfaces and protocols. Now, the avalanche of pervasive computing infrastructure pushes us to raise the level of plug-and-play to very heterogeneous devices and contexts. These problems cannot be solved only with appropriate protocols. Our in-place service infrastructure needs to understand what exactly the service needs to look like. Here is the place where Cognition and CVO are called to help at virtual level.

In the following we present some very simple examples.

Our homes and offices are now intensively populated with devices: surveillance cameras, printers, phones, computers, and a large variety of sensors such as the ones for light and temperature or humidity. We need to or want to share these resources, attach security to them, use and link them (controlled or not) with needs, intentions and processes. We need to give a meaning, real time effectiveness and for this fact we need to implement cognition capabilities. That is the paradigm of smart home.

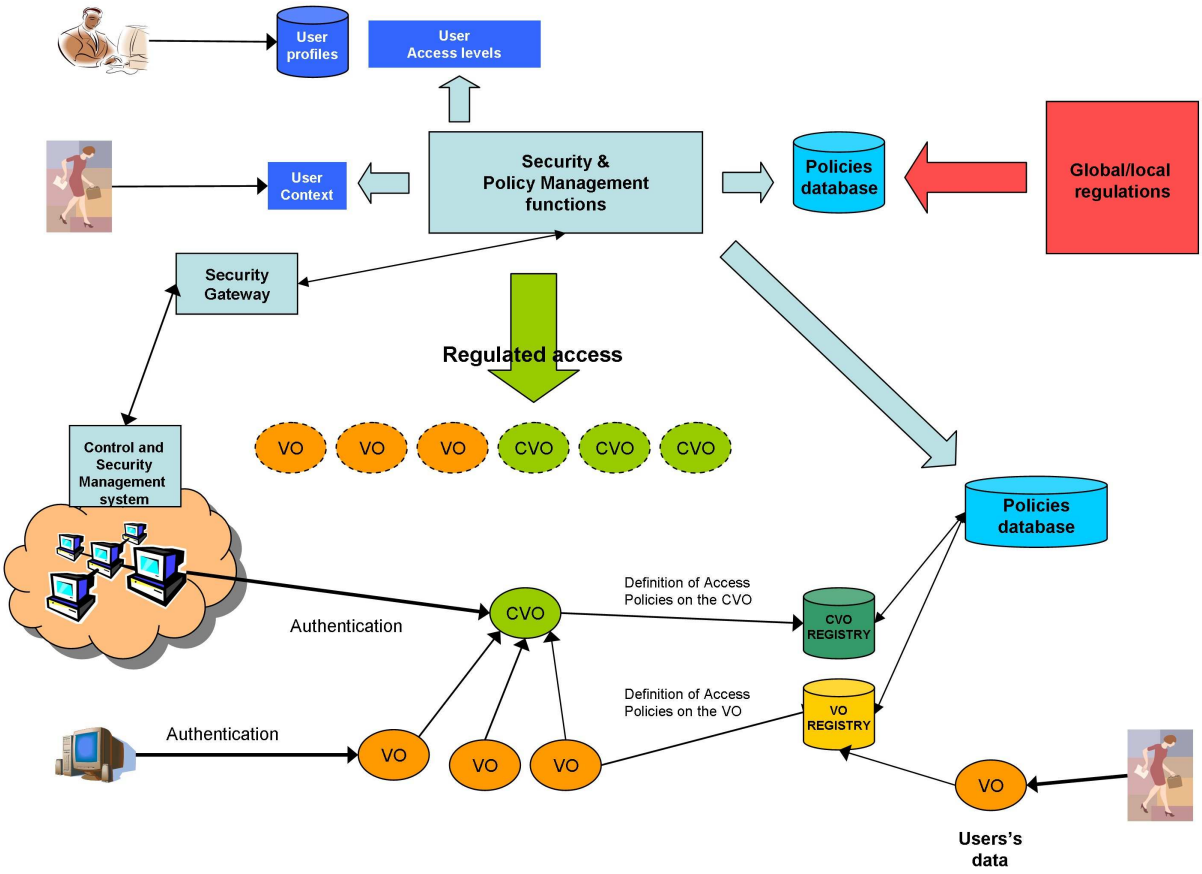

Fig. 3.1: Security framework and functions

Lets imagine that Mary is living with her family, including her father who needs permanent care due to a chronic health issue. Due to her daily duties Mary needs to travel in a neighboring city for some meetings. Her father remains at home, monitored by a number of specialized devices and some control cameras over the usual living space. Mary expects that in case of a possible critical situation, the problem is detected, relevant information is gathered, and specialized medical team and herself alerted in a consistent and timely manner (like a message on her smart phone with a short report attached). In a classical perspective, this kind of particular scenario requires specialized services developed, including spatial understanding of the problem (like the correlation of a body falling in a room with the variation of a body parameter).

Now lets have a look how the problem is tackled using VO/CVO and cognitive infrastructure. A person who needs medical monitoring is equipped with sensors specialized in various measurements for continuous supervision. All these sensors have their own virtual objects representations usable in order to trigger adequate reaction. The same principles apply to an accelerometer or a video camera (see the example of a body falling). Based on these primitive virtual objects, a cognition enabled infrastructure will be able to infer a CVO who refer all relevant virtual objects and compose the relevant services used to solve the case: calling emergency services, family being alerted, etc. The same can be expressed regarding the time dependence between all observed variations. As a consequence, a custom problem is expressed in terms of a set of cognitive enabled strategies and policies. Even more, the cognition doesnt remain fixed but ask for users/actors active feedback and incorporate successful executed steps and experiences. This scenario is not restricted just to a sensing approach, but can involve in the loop actuators. One benefit is the real time inclusion of correlated actuators effects in the cognition correction.

5. Conclusions. This paper described a new cognition based framework approach to manage the complexity, huge amount of data, systems and services which the Future Internet of Things will be comprised of. The approach is based on the concept of CVO and VO to simplify the management of heterogeneous systems and data. Specific features of digital and physical objects can be represented as VO attributes. Security framework and functions regulates the access to CVO/VO through a sticky management policy.

Future developments will investigate how to address scalability of the proposed approach: the proposed architecture must be able to support millions of CVO/VOs, make them accessible from various instances deployed in various domains and offer the necessary mix of autonomic and learning capabilties for optimal matching of demands with offer in an un-reliable, resource access, energy and communication constrained environment.

REFERENCES

[1] M. Uusitalo, *Global Vision for the Future Wireless World from the WWRF*, Vehicular Technology Magazine, IEEE , vol.1, no.2, pp.4-8 (2006)
[2] *IoT-A Website*, Available at http://www.iot-a.eu/
[3] *CASAGRAS2 Website*, Available at http://www.iot-casagras.org/
[4] M. Weiser, *The Computer for the Twenty-First Century*, Scientific American, pp. 94-10, (1991)
[5] J. Schonwalder, M. Fouquet, G. Rodosek, I. Hochstatter, *Future Internet = content + services + management*, IEEE Communications Magazine, , vol.47, no.7, pp.27-33, (2009).
[6] D.D. Clark, J. Wroclawski, K.R. Sollins, R. Braden, *Tussle in cyberspace: defining tomorrow's Internet*, Networking, IEEE/ACM Transactions on , vol.13, no.3, pp. 462- 475, (2005).
[7] S. Vinoski, *Web services interaction models. Current practice,* IEEE Internet Computing, , vol.6, no.3, pp.89-91, (2002).
[8] M.C. Mont, S. Pearson, P. Bramhall, *Towards accountable management of identity and privacy: sticky policies and enforceable tracing services,* Database and Expert Systems Applications, 2003. Proceedings. 14th International Workshop on , vol., no., pp. 377- 382, 1-5 Sept. 2003.
[9] F. Toutain, A. Bouabdallah, R. Zemek, C. Daloz, *Interpersonal context-aware communication services,* IEEE Communications Magazine, , vol.49, no.1, pp.68-74, January 2011.
[10] W.A. Alrodhan, C.J. Mitchell, *Enhancing user authentication in claim-based identity management,* Collaborative Technologies and Systems (CTS), 2010 International Symposium on , vol., no., pp.75-83, 17-21 May 2010
[11] J. Hoebeke, G. Holderbeke, I. Moerman, M. Iacobsson, V. Prasad,. N. I. Cempaka Wangi, I. Niemegeers, S. Heemstra de Groot, *Personal Network Federations*, Proceedings of the IST Summit 2006, Myconos, Greece, June 2006

Edited by: Marcin Paprzycki
Received: May 1, 2012
Accepted: June 15, 2012