

Extending the Internet of Things

Gilles PRIVAT
Orange Labs

Abstract: This article proposes a theoretical and practical extension of the IoT, taking in all "things" that can be sensed by sensors, without requiring them to be fitted with a tag or a digital network interface. These physical entities, whatever they may be (legacy appliances, passive items, subsets of physical space), become nodes of a broader network, extending the internet of sensor/actuator devices. We explain how such an evolution for environment-to-information interfaces draws upon a similar, long-standing evolution of human-to-information interfaces. Multisensor acquisition of physical context supports this extended IoT, bypassing the need for network-ready identification of target entities. We describe a three-layer reference architecture for an infrastructure supporting the integration of applications into the extended IoT. We show on a few examples how this can expand IoT applications and endow them with features of robustness, scalability and self-configurability.

Key words: Internet of Things, directed graph, physical context, multisensor data fusion, pattern recognition.

■ Introduction

Beyond the internet of devices

Most mainstream visions of the "Internet of Things" come down to extending the range of devices that may become connected to networks, moving from Wifi or cellular to RFID, Zigbee or their equivalents. The oft-repeated rationale is straightforward: there are trillions of such low-end devices out there, waiting to get connected, when billions of humans and their regular devices already are. If a new variant of Metcalfe's law ¹ applies, the promise of these "things to things" connections appears almost limitless.

¹ According to which the value of a network increases in proportion to the square of the number of nodes connected to it. A variant with lower exponent could possibly be derived for networks of non-peer entities, i.e. nodes without full configuration and bidirectional communication capabilities, as is the case for most "things" in the original or extended IoT.

Under such earlier catchphrases as "smart devices", "communicating objects", "pervasive networking" or M2M, it is no surprise that the telecom industry had been embracing this evolution as a legitimate extension of its territory, well before the "Internet of Things" gained currency as the new buzzword of choice.

Yet, compelling as it may seem, this "things-to-things" vision misses the crux of the evolution: this is not only a quantitative extension of existing person-to-person networks, it is a genuine quantum leap. By connecting "things" that are deeply embedded in the physical environment, ICT systems become strongly coupled with all kinds of physical systems. This opens up entire new domains that were so far entirely outside the reach of ICT, or for which information systems were disconnected from the corresponding physical plant/system/process, requiring manual configuration and manual data entry to couple the two.

Moreover, as viewed from the confines of the telecom industry, these early attempts at redefining communication beyond person-to-person have created some confusion because the distinction between the different categories of new things/devices that became attached to networks was not always made clear, especially when mobile phones and their avatars were added to the mix. Devices for which sensors and actuators are used exclusively to support regular human interfaces² should not be counted in for the IoT proper. The devices that make up the Internet of Things in its mainstream, yet strict sense, are networked sensors, actuators, devices equipped with sensors and/or actuators, or more generally networked devices that are not primarily ICT devices, equipped with physical interaction capabilities that correspond to their primary function. This Internet of devices comprises things/devices that are, in a proper sense, "embedded" in their physical environment, adding information processing and transmission capabilities to this environment³.

In this view, the outer border of the digital network is still the sensor or actuator itself, beyond which is the physical world. The revolution of pervasive networking that led to the multiplication of these connected sensors and actuators (PRIVAT, 2006) afforded an order-of-magnitude

² Classical IT interface devices and telecom terminals use sensors for information input from their human users and actuators for information rendering to their users: this is different from using sensors to capture data from the environment and actuators to modify this environment.

³ This may in fact include regular ICT devices such as smartphones if they are used for distributed network sensing rather than only information input/output from their users.

enlargement in bandwidth between networks and the physical world. Yet, for all their transformative role, this current generation of sensors and actuators do not correspond to the ultimate possible displacement of the network border. How this border may shift further to include real-world "things" is the next stage of the evolution that we intend to describe.

Drawing upon human interfaces

Most grand schemes devised for the classical supply chain management applications of the Internet of Things, such as the EPCglobal Network ⁴, the uidCenter ⁵, or "Internet 0" (GERSHENFELD, KRIKORIAN & COHEN, 2004) attach a universally unique, network-ready digital identity to these physical things, be it their General ID, ucode or IP address. This amounts to digitizing these "analog" things, or to shifting the border between the digital and physical worlds further towards the digital.

In the latest evolutions of human interfaces ⁶, this border has been moving in exactly the opposite direction, which amounts to making the digital/information/virtual world appear more like the analog/human/physical world. Human interfaces are designed so as not to force human users to meet the digital information world on its own terms, they should come closer to an interaction between humans than to an interaction between programs or networked entities. The entire agenda of ambient intelligence (STREITZ & PRIVAT, 2009) and such ideas as "perceptual interfaces" (REEVES & NASS, 2000) or "tangible interfaces" (ULLMER & ISHII, 2000) bear witness to this. The difference between data input through a keyboard and command line interface and input through a software personal assistant with voice recognition should make this clear. A less obvious example is the replacement of clicking on a menu item by the grasping of a tangible interface object that physically impersonates the same digital entity. Obviously, the interface is moving much further into the analog world in the latter case, and it requires more sophisticated sensing and perception capabilities on the part of the system.

⁴ <http://www.epcglobalinc.org>

⁵ <http://www.uidcenter.org>

⁶ This does not refer to virtual reality, which could be considered as a counter-trend (or a mere extrapolation of past trends).

These opposite evolutions have each been advocated for valid and widely accepted reasons in their own right. As for human interfaces, convenience and ease of use are not the only reasons for the un-digitization trend: robustness, graceful degradation and reliability could be complementary reasons, though this is not yet obvious with, e.g. the replacement of keyboard text input by error-prone speech recognition. It would be clearer if we could replace a password input by some 100% foolproof face or palm recognition software. For supply chain management applications that digitize everything from cattle to laundry detergent boxes, attaching universal identifiers appears mandated by the need to scale up the efficiency of digital data management to the physical world.

The main thesis of this article is that we can and should apply ideas drawn from the domain of human interfaces to networks of physical things: communication with and between these need not be forcefully digitized, it may retain the specific properties of the physical world in which these things belong, and gain new benefits from this. We explain in the following how sensor networks can be extended and consolidated into an infrastructure that supports un-digitized "thing-computer interfaces" inspired from ambient and context-aware human-computer interfaces.

■ Related work

A few European projects, among which IoT-A⁷ and FI-ware⁸ (DE, ELSALEH, BARNAGHIP & MEISSNER, 2012), propose a reference architecture for the Internet of Things that highlights the distinction between networked devices and real-world things (entities), in a sense close to the one proposed here. This distinction is matched to different root categories in the ontologies that support the relevant semantic models. These models are domain-specific for physical entities of the environment, and intended to be shared with applications. Our work goes further in extending the graph network model to these things and, crucially, the automatic discovery and configuration mechanisms that make it possible to integrate them in the network as if they were regular devices.

⁷ www.iot-a.eu

⁸ www.fi-ware.eu

Network and distributed software infrastructures with capabilities for spontaneous "zero-configuration" discovery and integration of new devices are a clear inspiration for the solution presented here. These infrastructures differ widely in the levels of interoperability they include in their scope. At the lowest end is the spontaneous configuration of network-layer addresses with mechanisms such as the automatic assignment of private ("link-local") IP addresses. At a higher level, distributed Service Oriented Architectures⁹ view devices through the description of the services they support. These services may get discovered and advertised once the devices are connected. A number of research and prototype solutions (SONG, CARDENAS & MASUOKA, 2010) have been proposed to extend these mechanisms to the semantic level. This would relax the requirement for prior definition of the corresponding services in the format required by the standard, for them to be discovered. Yet it does still require interoperability at the lower levels. However ambitious and high-level they may be, all these infrastructures are strictly limited to classical networked devices. They require that the target devices be natively endowed with network-ready interfaces or tags, and that the corresponding interfaces comply with required standards at all appropriate levels. The most constraining requirement is that devices need to be fully known beforehand as instances of extremely specific types¹⁰ for these mechanisms to work. Identifying the device to a general category (such as "printer" or "display") is very far from sufficient for a particular instance of this category to be integrated in a regular service-oriented architecture.

It is in these regards that the proposed approach differs most radically from existing infrastructures. It requires neither equipping devices with standard interfaces, nor pigeonholing them to a specific type. Devices can be identified by approximation to a very generic model, and the system should be able to integrate them on the basis of this minimal information.

Under the general "ambient intelligence" agenda (STREITZ & PRIVAT, 2009) the "smart space" research strand has targeted the physical environment as a primary basis for context grounding. In this view, devices are just transparent intermediaries and the environment itself may become an interface for human interaction. In placing the focus "beyond devices",

⁹ Such as UPnP (www.upnp.org) or DPWS (<http://docs.oasis-open.org/ws-dd/ns/dpws/2009/01>).

¹⁰ This type information includes at least to the manufacturer, the brand/make, the commercial name/model, the year of issue/version number.

this has much in common with the approach we propose here, which can be extended to relevant subsets of the environment as target physical entities. These "space entities" can be represented and modeled in a way very similar to "graspable" entities, using the same infrastructure to interface with applications that use high-level context information about the environment.

■ The extended IoT of sense-able/actionable things: reference model

Extending the network to "sense-able" things

We start from a narrow mainstream view of the Internet of Things as a network of sensor-equipped devices and examine how to make the external interfaces of this network more "thing-friendly", as if they were perceptual human interfaces. Instead of constraining things to adapt to the network, we make it possible for the network to adapt to them. Instead of enforcing their uniform digitization, we try to take them as what they are, analog physical things.

Let's for the time being take the example of a single sensor, e.g. a camera, and suppose we have a "thing" recognition and monitoring software analyzing the data acquired by this camera. Assuming this, we can consider that every individual thing or physical entity within the field of view of this camera becomes ipso facto a "networked thing", provided it can be recognized and monitored by this software coupled to the camera. This means it can have a presence on the network, without requiring an RFID tag or even a digital optical code (such as a 1D or 2D barcode) for this. We proposed in earlier work (PRIVAT, 2012) to extend the notion of "phenotropics"¹¹ as a broad conceptual basis for the use of thing recognition and monitoring through sensors as a new type of analog network interface. We will not be delving further into these theoretical aspects here, focusing instead on practical and implementation concerns.

In this view, the range of things that may become indirectly connected to the network can extend much further than sensor devices themselves, to all

¹¹ A concept and noun originally forged by VR pioneer Jaron Lanier, who viewed it first as a way to make the internal interfaces of information systems more robust and adaptable.

things that are individually identifiable by a sensor. This extension of networks to new nodes "beyond sensors" is represented in figure 1 by a set of directed "sensor to thing" links. It is not a routine incremental and quantitative extension, such as results from adding a new wire-line or radio-based protocol. It is a qualitative leap that makes it possible to integrate all discrete or bulk "stuff", as it is, without requiring any digital identity or network interface whatsoever, and without adhering to any kind of predefined standard, at any level, for this connection. There is no prior barrier to the integration of new things. It is by nature universal as it requires no prior standardization of any kind of digital code or interface.

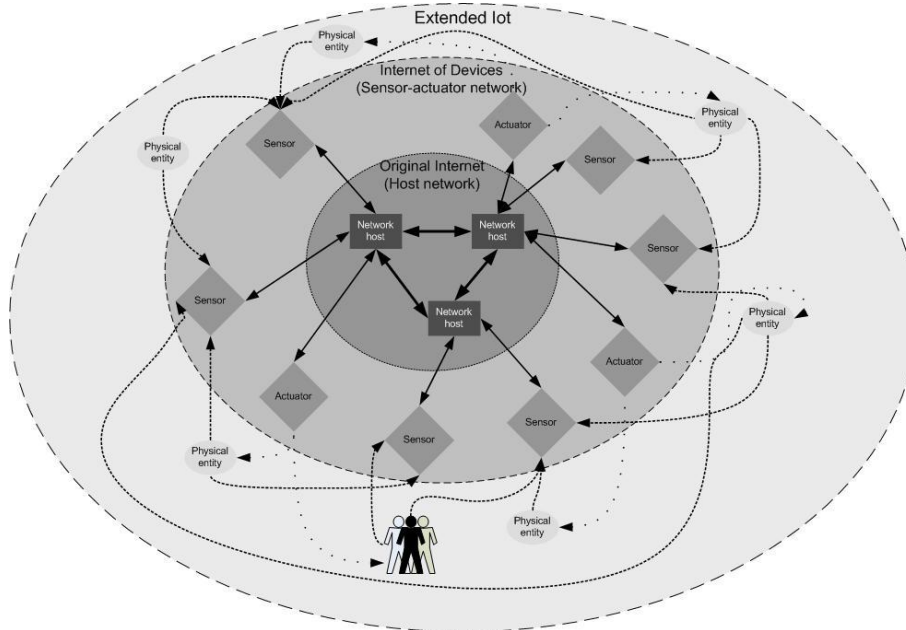
Extending the network to "actionable" things

As the counterparts of sensors, actuators transduce numerical variables into physical ones. They enact modifications of the physical environment and the effects of these are either sensed directly by sensors, or indirectly, through passive things which are modified by the actuators. These new physical "actuator-to-thing" links complement the "sensor-to-thing" links. Together they make up a directed graph, or virtual network that we have proposed to call a stigmergic network (PRIVAT, 2012), by reference to a biology-inspired concept of complex systems theory¹². This "actuation graph" is again overlaid upon the wire-line/wireless data network through which the corresponding sensors and actuators receive or transmit their respective numeric data (figure 1).

What we integrate in the network here is not new nodes, but new links that "close the loops" of sensor-actuator networks in a way that does not use the modalities of classical networks and complements them. New qualitative system-wide properties can be analyzed in this double coupling of sensors and actuators: through the network and through their common physical environment. This coupling may result in both extremely useful and potentially undesirable or harmful effects, making this a key challenge for future research. We will not elaborate here on these system-theoretic aspects.

¹² Stigmergy, a concept originally proposed by Pierre-Paul Grassé from his observations of social insects, refers to indirect communication between agents by the modification of a shared physical environment.

**Figure 1 - Internet of Things extending the Internet of Devices
(extension links as dotted arcs)**



Extended Internet of Things, or Web of Things?

If we called "Internet of Devices" the network of sensors and actuators (devices that are attached to networks in a traditional sense), should we call the extended network proposed above the "Internet of Things" proper, or the web of things? Much as the early World-Wide-Web that we know of was a virtual network of hyperlinked static HTML documents overlaid on top of the Internet, the network of sensed things can be considered as a virtual network extending an internet of devices and sensors, which is itself an extension of the early Internet¹³. As a virtual network (a graph in mathematical terms), this network of things comprises a far larger number of nodes than an IP network ever will, just as the web has many more documents than internet hosts. Another key difference is that, whereas the graph representing an IP network is non-directed¹⁴, the graphs representing

¹³ Represented by the inner circle of "hosts" in figure 1.

¹⁴ This means its links (*edges* in graph theory terms) are bi-directional.

either the classical document-centric web, or the extended Internet of Things, are *directed* ¹⁵.

Even though the so-called "web 2.0" has been more of a fuzzy amalgamation of hype than a distinct technological evolution, no less an authority than Vinton Cerf has claimed that the "web 3.0" would be the future Internet of Things. The "web of things" is an alternative name that could be used, hadn't this phrase already been proposed (GUINARD & TRIFA, 2009) to capture the use of lightweight application-level web-based protocols for the Internet of Things. To avoid confusion, we will keep to the phrase "extended Internet of Things" (using "xIoT" for short) in the rest of this article.

Contextual interfaces and multi-sensor pattern recognition

The notion of bilateral "sensor to thing" links presented above is a simple abstraction of the more diffuse, multilateral reality of context sensing that should actually apply for the identification and monitoring of things in the extended IoT. Here again, a very useful lesson can be drawn from the evolution of human interfaces.

In the course of moving away from the digital border of networks, human interfaces have become context-aware, which means they have been evolving from a simple bilateral human to device relationship to become mediated by the virtual and physical environment in which the interaction takes place. To focus here on physical context, a context-aware interface takes into account other sensor inputs than the primary, explicit user inputs, like when the user's location ¹⁶ is brought to bear to scope a request for some local service, without the user having to specify it explicitly. In a broader view, context-aware interfaces amount to using the entire physical environment, rather than one single device, as an interface. This environment becomes a smart, perceptual environment, where all sensors

¹⁵ This means its links (*arcs* in graph theory terms) are uni-directional. For the extended internet of things only the extension links (things → sensors and actuators → things arcs) are unidirectional, whereas all network links in the original web are unidirectional: this is related to the difference between an extension graph, which is a superset of the original graph, whereas an overlay graph like the web is in a different plane, its nodes being mapped to the nodes of the underlying graph.

¹⁶ Acquired through a sensor-equivalent location-determination technology which, whatever it is, can be considered as providing implicit secondary data complementing the user's explicit input.

are federated, acquiring low-level context data that is fused and interpreted to become high-level context.

A similar notion of physical context should apply to the "things to networks" interfaces, for them to become "less digital". All sensors of a given environment can be brought to bear in order to "connect" things. We again assume that we have available through the "Internet of Devices" a federation of distributed networked sensors that make up a "smart environment". These sensors are not dedicated to one application and they can all provide useful context data about this environment. Networked "things" in this indirect, extended sense may then comprise all "stuff" that can be sensed by pattern recognition software operating on top of these federated sensors *working together*, potentially overcoming their individual limitations as single-modality¹⁷ devices. Whether they are primary sensors, or provide only complementary data, the identification and monitoring of "things" in this environment will use them jointly, opportunistically, taking into account their data inasmuch as it is relevant.

The kind of pattern recognition used to identify and monitor things on this basis is very different from classical pattern recognition based on separate modalities, such as used in computer vision or speech recognition. Pattern recognition used in this way may rely on classical multi-sensor data fusion (HALL & LLINAS, 1997), but when using very basic sensors such as passive infrared, door contact or electrical sensors it is in fact much simpler than when dealing with rich and complex signals such as video or audio. It assumes primarily the detection of temporal coincidence of events from different sensors (as potentially coming from the same physical entity) and the application of simple filtering rules to these multi-sensor events. The consolidation of these events will then depend on the corresponding model of the originating physical entity. More concrete examples and descriptions of this are provided in the following sections.

¹⁷ Modality is used here in a sense derived from its use in human computer interaction, but is not limited to human sensory channels: a modality is a type of physical variable or phenomenon that is measured or detected by a sensor, such as temperature, pressure, location, movement, etc.

From universally unique to contextual identification

Things that are connected to the extended IoT in the sense proposed here need not have an explicit, pre-assigned and pre-registered universally unique identification attached to them, as is normally a prerequisite for mainstream IoT "things" with classical technologies such as RFID. It is in fact sufficient for most applications to assume that things are *implicitly* identified. This means that they are identified on a relative, non-universal basis *in a given context* or scope. It is the knowledge of the context, combined with the local identification, which can make the identification of things global and universally unique if needed. The software proxy (representative) of these things in the network will maintain this identification in a way that can be used by applications. Human computer interfaces are again evolving in the same direction, requiring a unique universal identification from users only when it is needed, as most users will prefer a contextual identification that preserves their privacy (such as e.g. a session cookie for web-based applications) when it suffices.

The key difference with traditional network-based identification is similar to what was already mentioned for network connection. Contextual identification does not require prior standardization and shared knowledge of a set of codes or protocols for exchanging this identification. It does not require either a prior registration of the object in a database. The entity may be identified on the basis of its physical features, as these are sensed by the available sensors, associated with required context data as a complement. The "recognition" of the entity and its association with a known category relies on generic, publicly available knowledge, not on a matching with some proprietary database. If not unique as an instance of this category, context (such as location) comes in to complement the category-based identification¹⁸.

■ Extended IoT: reference architecture

From the proposed conceptual model, we derive a layered software/network architecture that can serve as a reference framework for

¹⁸ Universally unique identification à la EPC global may of course still be a requirement for some applications such as supply chain management, but this does not mean it should be enforced when it is not needed.

the implementation of ICT systems based on the extended IoT as defined in this article. This architecture draws an analogy to the layering of both standard networking models¹⁹ and computer architectures, where the underlying physical hardware is abstracted away, hidden under successively more abstract and hardware-independent interfaces provided to applications (figure 2).

Extended IoT framework

From the viewpoint of an application using the extended IoT, the problem addressed can be stated in the most general possible way as follows: an ICT system is set up to acquire data from a large-scale physical system and, if need be, control it in return. Shared sensors and actuators are distributed as monitoring and control points through this physical system. A set of entities, subsystems of the overall physical system, are defined as the components of the overall physical system that are relevant for being controlled and monitored by the targeted application. These subsystems are distinct physical entities. They are fully-fledged physical systems in their own right: they will be the nodes of the extended IoT for this application. The sensors and actuators are not target entities themselves. They are used just as transparent intermediaries.

The ICT system will "shadow" each of these xIoT nodes individually through matching software components (proxies) that will offer to applications the required interfaces to the extended IoT in this environment. The ICT system should have the capability to create and configure these components automatically. This is required for the initial configuration stage and when a change in the environment triggers a reconfiguration. The configuration includes the automatic association with the entity proxy of the interfaces to the subset of sensors and actuators that are used as intermediaries for the monitoring and control of a given entity.

¹⁹ Such as the (7 layer) ISO or (4 layer) internet models.

Abstraction of xIoT nodes

We choose to represent the target entities (xIoT nodes) through simple discrete-state models transitioning on discrete events²⁰. Their digitized states, possibly complemented with relevant continuous-valued attributes, are then stored as the state of the proxy of the xIoT node. These models represent a simple yet adequate common denominator abstraction of reality for many practical xIoT use cases. They provide, as the centre points of discrete classification clusters, the necessary "anchor" for making sense of multidimensional sensor data without resorting to complex pattern recognition techniques.

Self-configuration and reconfiguration

Self-configuration makes it possible to identify and integrate spontaneously and automatically new entities into the xIoT network. The process goes through the following stages:

- Detection of meaningful sensor event as corresponding to a new entity
- Creation of a generic entity proxy in the corresponding entity abstraction layer
- Association of this entity proxy with the sensors that provide data about this entity
- Update of the entity model from additional sensor data
- Assignment to more specific entity category and more specific entity model
- Association with complementary sensors

From step 6, the configuration process may iterate in a loop including steps 4 to 6 until the most specific available model is reached.

Re-configuration refers to the continually on-going adjustment of the system to account for changes in the environment, such as removing or

²⁰ This means their status is captured by a state vector which is a discrete-valued function of time with discrete asynchronous (event-based) transitions, defined in a system-theoretic sense as encapsulating the necessary and sufficient information to obtain the next states and outputs of the system given its next inputs. Cf. CASSANDRAS & LAFORTUNE (2008) for an introduction to discrete event systems.

adding a new sensor/actuator, moving an entity or removing it. It is triggered by a mismatch between the sensor data and the entity model. From there, it involves a backward traversal of the graph of entity models ²¹ until the data matches again with the model. The configuration may then start over, just as an initial configuration, going through the very same stages.

Entity group representation

Proxies representing individual physical entities are all distinct software components at the same hierarchical level and do not contain one another, even if the corresponding entities have such containment relationships: for example, the proxy of a room will not contain the proxy of an appliance, even if this appliance is inside the room. An additional separate layer of "entity groups" is needed to represent such aggregation or containment relationships between physical entities, with a 1-to-n or n-to-1 mapping to the physical entity layer. A virtual entity may thus link to several physical entities or a physical entity may link to several virtual entities. The relationship between these different layers is represented in figure 2.

Interface to applications

The interfaces that are exposed to applications from the proposed xIoT architecture abstract away sensor and actuator data at a level corresponding to the states and associated attributes of the entities, as defined above.

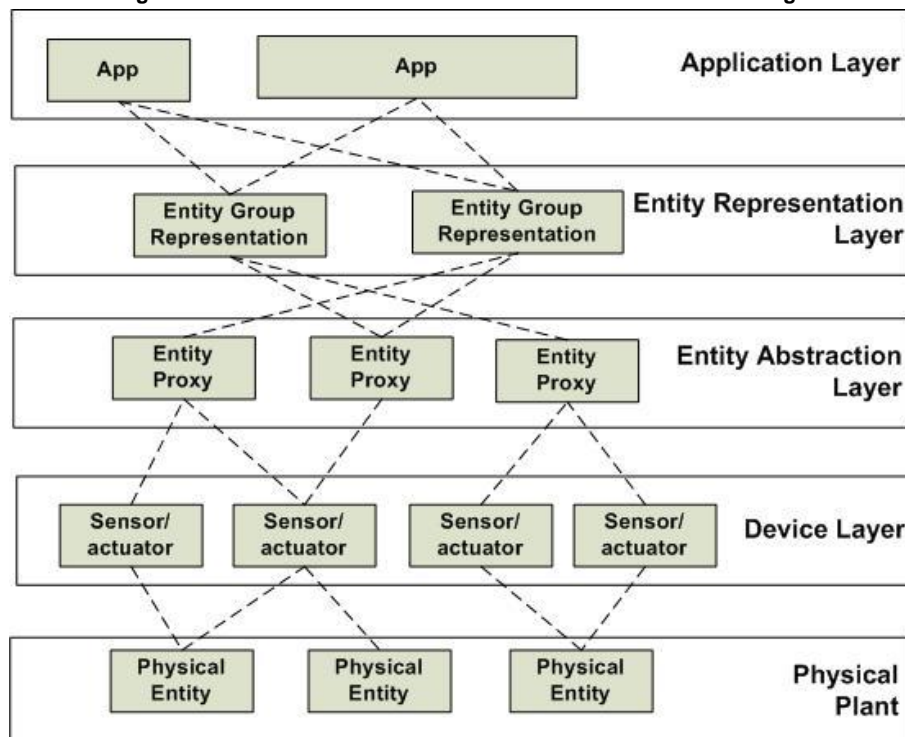
For monitoring an entity through the xIoT, an application can obtain the instantaneous state of this entity as the discrete state of the corresponding entity proxy, associated, if required, with complementary continuous-valued attributes. This discrete state is estimated as a result of the fusion, aggregation, consolidation and classification of data from sensors associated with the entity.

For control purposes, an application can effect a change in the state of an entity to another admissible state through the entity proxy that relays this

²¹ These models and their inheritance relationships make up a directed acyclic graph (DAG), which results from merging the arborescences corresponding to different complementary classification criteria. This DAG is traversed from the roots (the parentless; most generic models) to the leaves (the childless, most specific models) in the initial configuration phase, and traversed back from the leaves to the root when a reconfiguration is triggered.

high-level state-change control order to low-level control data for the associated actuators.

Figure 2 - Reference architecture for the extended Internet of Things



■ Application examples

Integration of legacy appliances in a home network

In the home environment, target entities are those that are relevant for being monitored and controlled by applications such as energy management, security/safety management or home automation, extending the home area network beyond its regular perimeter of ICT devices and state of the art home automation devices.

If we take home energy management as an example, examples of the target physical entities would be:

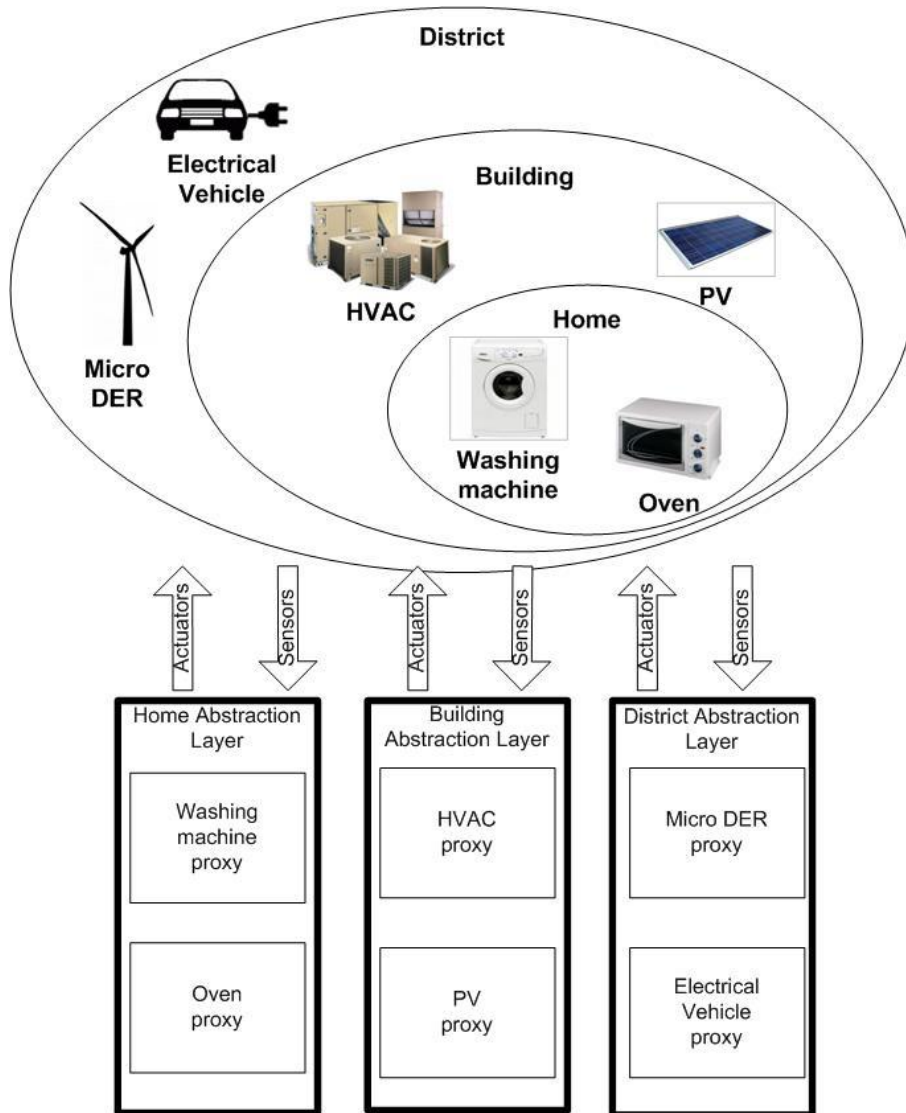
- appliances and devices of all types, including all pieces of legacy home equipment,
- rooms of the home,
- energy-relevant components of the home such as walls, windows.

For other applications, this might be extended to pieces of furniture, pets, or the home occupants themselves. These mostly non-digital entities have to be integrated in the home xIoT network in a way similar to what is done with regular networked entities. This means they have to be identified and matched to an existing model that can be specific or generic, exact or approximate. State of the art devices would afford this integration through a high-level SOA-like interface, but until they are widespread in the home domain (which will take a long time because of the slow turnover of home appliances), we still have to deal with legacy appliances whose only available interface is that of their mains connection. This interface makes it possible to identify these appliances through the characteristic features of the patterns exhibited by their electric power consumption through an electric power sensor (like e.g. an oven showing a steady plateau pattern whereas a washing machine has characteristic peaks and troughs). This electric current sensor will be the main sensor for mains-connected appliances, possibly complemented by other sensors available in the environment. When these appliances are identified and enrolled into the extended home IoT in this way, it becomes possible to monitor and control them as specific or semi-generic entities, even though this control is limited to the mains interface. This is not equivalent to what can be done through a state of the art data network (which would in principle make it possible to remotely program the appliance, or at least change its mode of operation) but it may still be sufficient for the purpose of monitoring and controlling it for energy management (HU & PRIVAT, 2011).

Multiscale energy management in the Smart Grid

The smart grid can be considered as the result of adding an IoT layer on top of the electrical grid. The smaller scales of the smart grid may involve the decentralized management of semi-autonomous units such as home, building or district "micro grids", where all kinds of electrical equipment and energy-relevant physical entities can get integrated in a local energy management system.

Figure 3 - Extended IoT supporting multi-scale decentralized energy management in the smart grid



As explained before for the home and building domains, these entities are widely diverse and heterogeneous, adding to the mix of home/building appliances the type of power equipment that gets connected to the distribution network. This may be classical electrical engineering machinery such as inverters and transformers, or the newer type for which the smart grid is precisely intended, like distributed renewable energy resources (e.g.

wind turbines or PV panels). We advocate that the "extended IoT" as proposed and described here is the proper approach to the design of the ICT layer of the smart grid because, among other reasons, it is essential to the viability of this extended and decentralized smart grid approach that the integration of these entities does not require manual configuration. The nested scales corresponding to home, building and district energy management systems, together with examples of the corresponding entities, are illustrated in figure 3 (HU & PRIVAT, 2012).

■ Conclusion

The proposed approach has been developed and validated so far in the home domain, on the basis a complete range of home entities, from electrical appliances to rooms. The home area network thus extended may be considered, if used for energy management, as a home-scale smart grid, intended to nest within larger scales of the smart grid using the same extended IoT approach. The larger scales that we may deal with will correspond to cities or city districts, for which smart grids are but one application. Relevant target entities for smart cities might be as diverse as lampposts, garbage containers, cars, pedestrians and building themselves. Besides this, all classical applications of the Internet of Things and smart environments could be revisited by applying this approach, relaxing the requirements for state of the art network interfacing and digital identification of physical entities, making them more open, "things-friendly" and ultimately, we hope, more widely accessible and successful.

References

- CASSANDRAS, C. G. & LAFORTUNE, S. (2008): *Introduction to Discrete Event Systems*, Springer.
- DE S, ELSALEH T., BARNAGHIP, P. & MEISSNER S. (2012): "An Internet of Things Platform for Real-World and Digital Objects", *Scalable Computing: Practice and Experience*, Vol. 13, no. 1, pp. 45-57. <http://www.scpe.org>
- GERSHENFELD, N., KRIKORIAN, D. & COHEN, R. (2004): "The Internet of Things", *Scientific American*, October.
- GUINARD, D. & TRIFA V. (2009): "Towards the web of things: Web mash-ups for embedded devices", in *Proceedings of WWW conference, Madrid, Spain*
- HALL, D. L. & LLINAS, J (1997): "An introduction to multisensor data fusion," *Proceedings of the IEEE* , vol. 85, no. 1, pp. 6-23, January.
- HU, Z & PRIVAT, G.:
- (2011): "Iterative Model-based Identification of Building Components and Appliances by Means of Sensor-Actuator Networks", *Computers in Buildings 2011*, Workshop on data models for energy-efficient buildings, Sophia Antipolis, October 2012.
 - (2012) "Representation and Self-Configuration of Physical Entities in Extended Smart Grid Perimeter", *IEEE Innovative Smart Grid Technologies (ISGT)*, Berlin, October.
- PRIVAT, G.:
- (2006): "From Smart Devices to Ambient Communication", Invited conference, *From RFID to the internet of Things*, EU Workshop Brussels, March 3 2006. ftp://ftp.cordis.europa.eu/pub/ist/docs/ka4/au_conf670306_privat_en.pdf
 - (2012): "Phenotropic and Stigmergic webs, the new reach of networks", *Universal Access in the Information Society*, Vol. 11, no. 3, pp. 323-335.
- REEVES, B. & NASS, C. (2000): "Perceptual user interfaces: perceptual bandwidth", *Commun. ACM* 43, 3 (March), 65-70.
- SONG, Z., CARDENAS, A. & MASUOKA R. (2010): "Semantic Middleware for the Internet of Things", *Internet of Things 2010*.
- STREITZ, N. & PRIVAT G. (2009): "Ambient Intelligence", Chapter 60 of *Universal Access Handbook*, Constantine Stephanidis, Editor, CRC Press. <http://www.crcpress.com/product/isbn/9780805862805>
- ULLMER, B. & ISHII, H. (2000): "Emerging Frameworks for Tangible user interfaces", *IBM Systems Journal* 39(3&4): 915-931.