Smart Building Energy Management : A "Future Internet" Perspective

Gilles Privat[1], Menelaos Perdikeas[2]
[1]*Orange Labs, Grenoble, France*
gilles.privat@orange.com
[2]*Synelixis Solutions, Chalkida, Greece*
mperdikeas@synelixis.com

Abstract

This paper reports on the work undertaken in the Finseny (Future Internet for Smart Energy) project on Smart Buildings as both domains of the Smart Grid and small-scale Future Internet platforms. A comprehensive set of use cases for building energy management have been identified, described here through high-level uses cases that are transversal to building categories. A. subset of key ICT requirements elicited on this basis are described. Going beyond present-day vertically integrated solutions, energy management for Smart Buildings needs to monitor and control all energy-relevant building subsystems, appliances and other physical entities in a non-ad hoc way, operating on top of a shared platform, a building "operating system" that can be provided to all building applications, abstracting away these entities behind the common interface of a generic driver-like model, mirroring the architecture of the Future Internet platform.

Keywords: Building automation**,** Building Energy Management System, Smart Space, Smart Grid, Micro Grid

1. Introduction

The importance of the building sector in the energy landscape can hardly be overstated. Buildings account for, overall, up to~70% of electricity consumption, ~40% of primary energy consumption and ~30% of GHG emissions[1].
Buildings are also characterized by a prevalence of legacy plant that places strong constraints on energy management solutions: it would be utterly pointless to design clean slate solutions that would work only for new buildings, or would require extensive upgrade.
The building sector is also highly diverse and heterogeneous, not only between building categories (such as residential, office, industrial, etc.), but also within these categories, which make the derivation of common denominator solutions all the more challenging.
Yet the building sector does present opportunities for energy management because it is, compared to other smart grid "domains" (especially the distribution network), where state of the art ICT is already most widely deployed.
In spite of its wide diversity of plant, actors and stakeholders, the building sector needs to derive common denominator ICT platforms, not only for energy management, but also for other building automation applications, simply because bespoke solutions are not economically viable.

---

[1] These figures do obviously vary in time and by country/region; they are given as an indication

We describe in the following the first steps of the task that has been tackled by the FINSENY project, towards deriving this platform on the same premise of scalability, reusability and openness as the Future Internet platform at large.

2. Defining target building domains

General assumptions

In keeping with received system design methodology, we started by defining a broad scope for this work, narrowing it progressively as we went along. We intended to address all types of buildings in a comprehensive way, as self-contained systems that encompass all the fixed, movable and mobile physical components of the building.
The information systems that manage general and energy-related functions of the building are not considered to be an integral part of the building system for delineating the target building as a physical system. They are considered to operate on a different plane, addressed separately under the ICT requirements phase. This means that use cases that are directly related to the ICT systems of the building, such as using multimedia communication systems for their own sake, or configuring, personalization and management of all ICT systems, are not addressed at this stage. On the other hand, use cases that correspond to specific "intrinsic" functions of the building but are in some way or another partially supported or assisted by ICT systems are taken in consideration and will translate into ICT requirements, along with those use cases that are not currently ICT-supported but are intended become so in the framework of the project.
The energy use of the building is assumed to comprise all potential local sources of energy, with emphasis on renewable sources, and all potential means of storage of energy, with emphasis on electricity. As an *energy carrier* for external sources of energy, *we restrict ourselves to electricity,* an assumption that is share across the FINSENY project.

Articulation with microgrids and distribution networks

Buildings of all types discussed in the following are supposed to be integrated either in microgrids, within which they are supposed to be "peers" at the same level as other entities connected to the microgrid (such as renewable energy sources), or directly to a distribution network. In both cases a pivotal 2 way interface to the microgrid or the distribution network implements a proper "separation of concerns" between these nested levels of system integration. This interface is a double 2-way interface, coupling information and power in both directions, from the grid to the building (downstream control information & power consumed from the grid by the building) and from the building to the grid (upstream status data & locally stored or generated power fed by the building to the grid).
As for control and data, this interface implements separation of concerns in a way that is merely conformant to received methodologies for the design of large information systems. More precisely this corresponds to the idea that the grid/microgrid should not have to know the details of the individual appliances and pieces of equipment (examples listed below) handled at the building level, only aggregate information being exchanged through the interface. If e.g. a load shifting or load shedding demand management request originating from the grid is transmitted through this interface, it need not and should not specify which appliance should be shifted or shed, it should specify only generic constraints (amount and duration of power to be shed) and it will be up to the building management system to decide which appliance should be shifted or shed, because only the building management system has the proper local context information to take a fully informed decision about this.

Intensional definition of entities included (equipment, appliances, components) in the target system

In a very broad view, the target systems comprises all parts of the buildings and all pieces of building equipment that have a direct or indirect impact on the energy input and output of the building. This includes all appliances/apparatuses that consume, generate or store energy, the components of the building such as walls and windows that regulate the exchange of energy between the inside and the outside, but also, in a more indirect way, subsets of the building such as floors or rooms and that make sense as separate units for managing energy in the building. Note that human users of the building are included either, depending on their role, as actors or as part of the internal environment and never as part of the building system itself

Buildings sub-domains targeted

For practical reasons, we limit ourselves to the following categories of buildings that correspond to specific interests of Finseny project partners

- Homes domain
- Residential buildings
- Office and public buildings
- Industrial building, esp. data centres
- Commercial buildings esp. hotels

3. Building use cases

Actors

Other domain-specific actor categories will be mentioned under each of the sub-domains addressed in the rest of the paper

Environment

In a comprehensive definition, the environment of the building comprises everything that has an influence on the state of the building excluding all specific other actors listed below and everything that is part of the system as defined above. Environment t can be specialized into external an internal environments., where the external environment comprises potentially some more specialized actors like neighbouring buildings or the weather, whereas the internal environment comprises internal factors that are directly under control of building systems, and are not within the perimeter of the building itself as defined above, such as activities of building users, inasmuch as they are not intentional interactions towards other use cases, or if they have side effects that impinge on the internal state of the building., or any factor of the external environment that “leaks” inside the building and cannot be controlled, e.g. natural lighting if openings cannot be controlled.

People and organizations

This is where the different categories of buildings defined below make the largest difference, so that we prefer to define these categories of actors separately for each of the building types as defined below.
Energy Service Companies take over building energy management entirely on behalf of the building owner or tenant, already play a well-recognized role for the management of non-residential buildings, but their role is so far limited in the case of residential buildings due to the lack of an established business model for them in this domain, by contrast again to office and public buildings where energy performance contracts are well-established.
Facilities managers play a role for all types of buildings except individual homes

External systems and physical entities

As for energy supply systems, the emphasis is, as per assumptions above, placed on the electricity distribution grid and microgrids in which the building is integrated, if relevant, but other energy carriers such as district heat or hydrogen could be taken into account. We normally do not consider the case of bilateral exchange of energy between two buildings, as it should be subsumed by microgrids.

Cross-domain high-level use cases

Monitor & control (manually) energy use

This use case corresponds to the most widespread current application of energy management, across all building categories. It may be considered as a stepping stone to more advanced optimization use cases, but it is also an integral part of any energy management system, keeping users in the loop, whether they are the end users of the building ‘as is the case in residential buildings, or facilities managers for office and industrial buildings

Optimize energy with global constraints, manage load/demand

In this use case, the building as a whole responds to constraints originating from the grid, mostly corresponding to demand shaping/load management, but also to generation if available. This is not equivalent to load management schemes as they existed in pre-smart grids, where the biggest individual loads inside the building or the home (e.g. water heaters) might have been controlled separately. In this use case, the interface offered by the building energy management system to the distribution network implements a proper “separation of concerns” between different levels of system integration : the external grid (or microgrid) does not “see” the individual loads of the building, it is the responsibility of the building energy management system to pass on load management constraints to the lower level, taking into account its own criteria, that are not known to the upper level, for this.

Optimize energy locally

This use case is complementary and always associated to the previous one, as local (building-scale)optimization is always performed jointly with global (grid-scale) optimization. It is emphasized here because it is not yet widely implemented or even understood. Building scale optimization may take into account criteria that are purely local to the building, such as the presence or activity of its occupants , to optimize energy efficiency in a strict sense, i.e. providing the same energy service with reduced energy input. Another dimension of local optimization is the local balancing of sources, loads and storage locally, which amounts to deal with the building as if it were a microgrid in islanding mode. Even if such islanding is rarely if ever possible, a more general criterion of global optimization can be to use locally generated energy as much as possible, which amounts to optimize the local loads with regard to availability of local sources.

Generate & store energy locally

The use of distributed energy resources and distributed grid storage is currently dealt with mostly at the microgrid and distribution network levels. This requires that these resources have their own direct interface to the upper scales, which is reasonable for the most important of these, but would not be optimal for proper complexity management if a number of such resources would exist inside buildings. For such smaller generation and storage resources, we propose to deal with them as a sub use case of local optimization as explained before, which means that these sources and storage means will be used primarily for local optimization and only secondarily as microgrid or grid-scale resources, and in the latter case they will be handled indirectly, aggregated through the building energy management interface

The following diagram represents the relationships between the use cases and actors in the case of homes. Similar diagrams have been derived for other building categories (1)

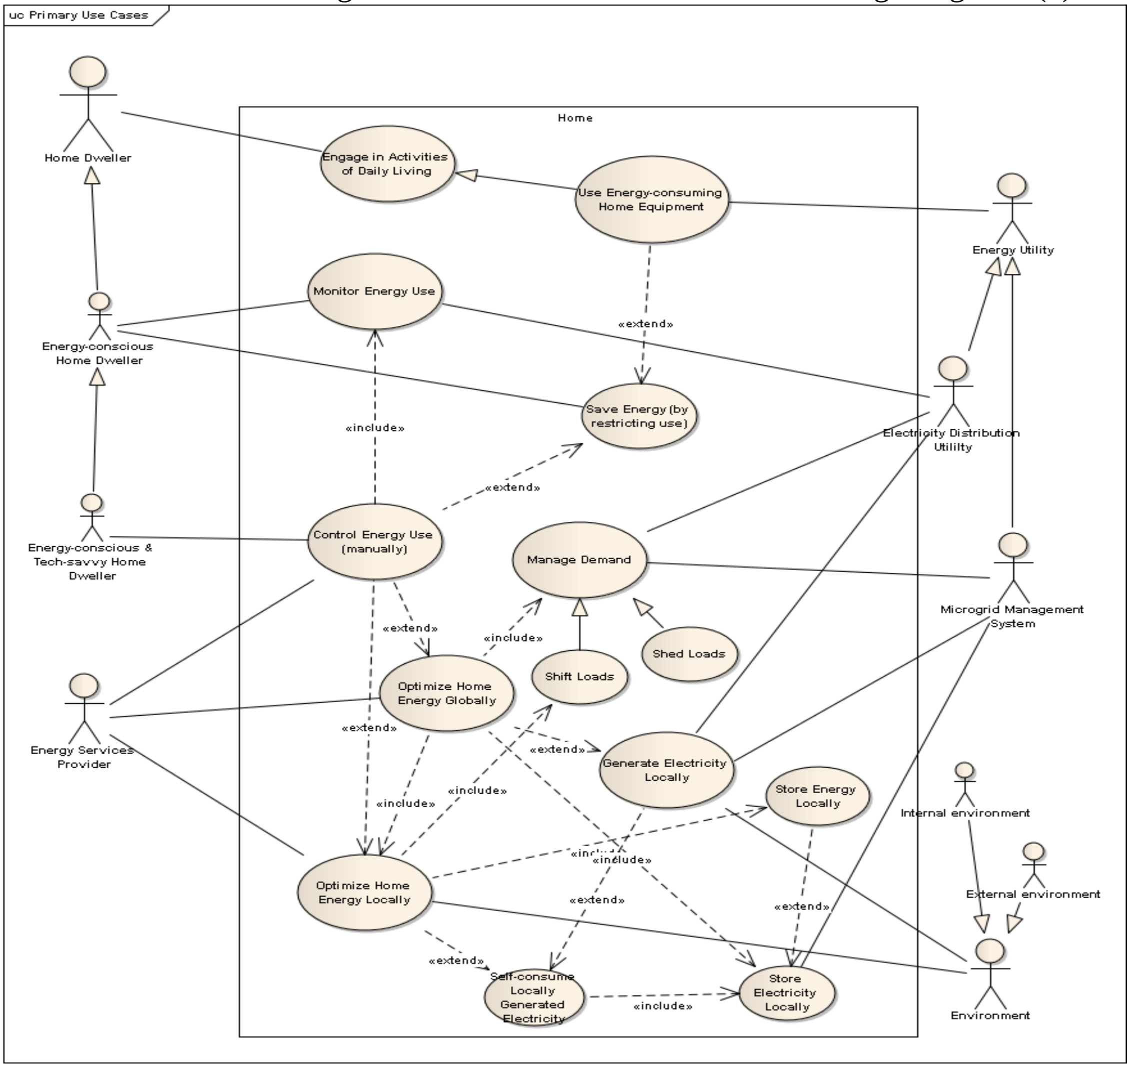

4. Buildings ICT requirements

ICT requirements are destined to serve as input for the Future Internet Core Platform project. A total in excess of 50 requirements were elicited for the building domain. We place here the emphasis on key requirements that are the least easy to meet or that challenge the state of the art.

Self-configuration

Most advanced applications of smart energy management in buildings require, and will increasingly require, automated configuration or self-configuration. Self-configuration could entail: (a) service discovery, (b) matching of appliances power profiles / signatures to a library of appliances, together with a domain information model similar to IEC's Common Information Model (CIM), to enable legacy devices to join the ecosystem (c) automatic recognition of appliances added or removed, (d) firmware upgrades and software patches. Requiring the residential user to manually configure appliances, download the latest software or patches and keep maintaining this infrastructure as appliances are added, removed, upgraded will seriously hamper end user acceptance. Historically, users have begrudgingly endured all sorts of glitches and software problems on their PCs and usually accept the time cost required to trouble-shoot common software or hardware issues.

However, this patience sufferance has been the result of: (a) reduced expectations (since all high-end technology is always assumed to be unreliable), (b) no alternatives and (c) the non-critical role played by the home PC (Internet browsing, entertainment hub). None of these conditions apply in the case of the appliances and ICT infrastructure to be deployed in smart buildings. In contrast to computer equipment: (a) people take for granted that a washing machine "just plays", (b) alternatives exist (revert back to the "dumb" version of the appliance) and (c) the smooth operation of an array of home appliances is critical, in some cases even life-critical. In end-user acceptance studies concerns have already been voiced about the complexity and the added clutter (for instance, in terms of wires or additional boxes) that these new "smart" technologies might bring in the home environment. Given such low tolerance and the fact that time is of essence in modern life, users will likely not tolerate too many frustrating experiences (such as having to figure out why the smart home energy management application doesn't "see" the new washing) before they abandon the "smart home technologies".
Most software / hardware / communications issues one experiences in daily life are usually are configuration problems or software bugs. Hardware problems are less common due to the less stateful nature of hardware and, hence, the possibility of more thorough testing. State results in the combinatorial explosion of different configurations and thus, the inability of thorough testing. Drawing upon lessons learned in the software industry the last twenty years we can identify the following self-configuration challenges for smart buildings:

1. Definition of a Plug-n-play architecture for building appliances.
2. Ability to automatically download configuration information (e.g. programs information, firmware updates, software patches) without requiring user supervision and without requiring a "reboot"
3. An always-on Internet connection to download the above
4. Sandbox and security model to ensure that malicious code cannot take control of the appliances and that infections are localized
5. Tools to automatically diagnose viruses (runtime protection)

Given the very high standard of reliability that this new breed of appliances and smart home applications will have to live up to, it may even be necessary to adopt development processes, tools and languages from the automotive / train / aviation industries instead of the computer industry. E.g. use of formal methods, mathematically-provable specification languages, model-driven code generation and / or functional languages as opposed to conventional object oriented or procedural languages whose stateful nature makes them impossible to exhaustively debug.

Semantic interoperability

Smart Buildings middleware should expose the functionality of appliances and electricity network resources, including legacy ones, in a semantically-rich way, not just in the form of a communication protocol but also at a higher level of abstraction and using concepts and models targeting the building domain, available in published ontologies. Appliances, apartments and whole buildings should publish, at their own level, standard open interfaces for reporting, monitoring and control of their energy-relevant physical subsystems for an ESCO, 3rd party service provider or Home Energy Management application software to use. Such interfaces could be provided in the Representational State Transfer (REST) style that is fast gaining traction in M2M systems (2). The underlying platform can also support other communication modalities like asynchronous messaging, publish / subscribe schemes and transactions, in addition to horizontal services like naming, persistence, etc. The rationale behind this requirement is to enable the development of novel software services for the building domain and also to reduce the costs of developing such software and increase its robustness. I.e. by being able to rely on well-tested facilities and services that are provided as part of the infrastructure, ad-hoc 3rd party software will be simpler and thus more predictable and stable, with fewer bugs, and easier to maintain and evolve.

Contextual security

The infusion of the home environment with remotely monitored and controlled smart appliances, ICT technologies, together with always-on Internet connectivity, will give rise to privacy and security concerns. Authentication and authorization technologies and solutions are well developed and understood in the ICT domain but their application in the home field poses new challenges: (a) hassle-free, even automatic, deployment and administration of security and authentication solutions, (b) security and encryption of data that can be used to monitor a home or profile its inhabitants whether such is held locally or remotely (c) taking advantage of physical proximity or known appliances / modes of communications to lower the authentication requirements (d) recognizing the criticality of different kinds of instructions and applying the necessary level of authentication to each. An example of (c) is that if an order to raise the house temperature is issued from a mobile device that belongs to the home owner no further authentication is necessary. Likewise if an order is issued from a computer connected to the home network or from a device that is reliably established to be in the premises. An example of (d) is that an order to dim the lights probably doesn't require any kind of privileges at all. Access control to the various applications and smart energy software subsystems inside a home should obviously be Single-Sign-On based. Typically computer network authentication protocols like Kerberos require non-trivial setup and administration effort; a solution for the home or smart building domain should be more easily deployed and maintained. Contextual security can be coupled with biometric authentication (e.g. face recognition from cameras that are perhaps also deployed as part of a security application) or even behaviometric authentication (e.g. based on gestures, typing rhythm or gait). This points to an integration model whereby applications expose not just black-box external interfaces but also internal features and services that are incidental or subordinate to their main use case (e.g. a surveillance application exposing a camera live feed that's used by an authorization application). Again, this is technology that is already in existence but the challenge is in deciding how, at what level, and where to integrate/deploy it within a smart building as well as how to exploit synergies with other applications and software stacks of that ecosystem.

5. Generic building-scale "Future Internet" platform

It is a pivotal tenet of the Future Internet programme as a whole that a set of common-denominator enablers should make up a shared software platform, a foundation for an internet that will be much more than a network, providing transversal services to applications in such areas the management of physical "things" and context awareness. The present situation in building automation is very far from this, as each specialized system (such as HVAC or security management) may be entirely closed and vertically integrated with its own networking and its own sensors and actuators. Going beyond these present-day siloed solutions**,** energy management for Smart Buildings should make it possible to monitor and control all energy-relevant building subsystems, appliances and other physical entities in a non-ad hoc way, operating on top of a shared platform à la Future Internet, a building "operating system" as pictured the diagram below.

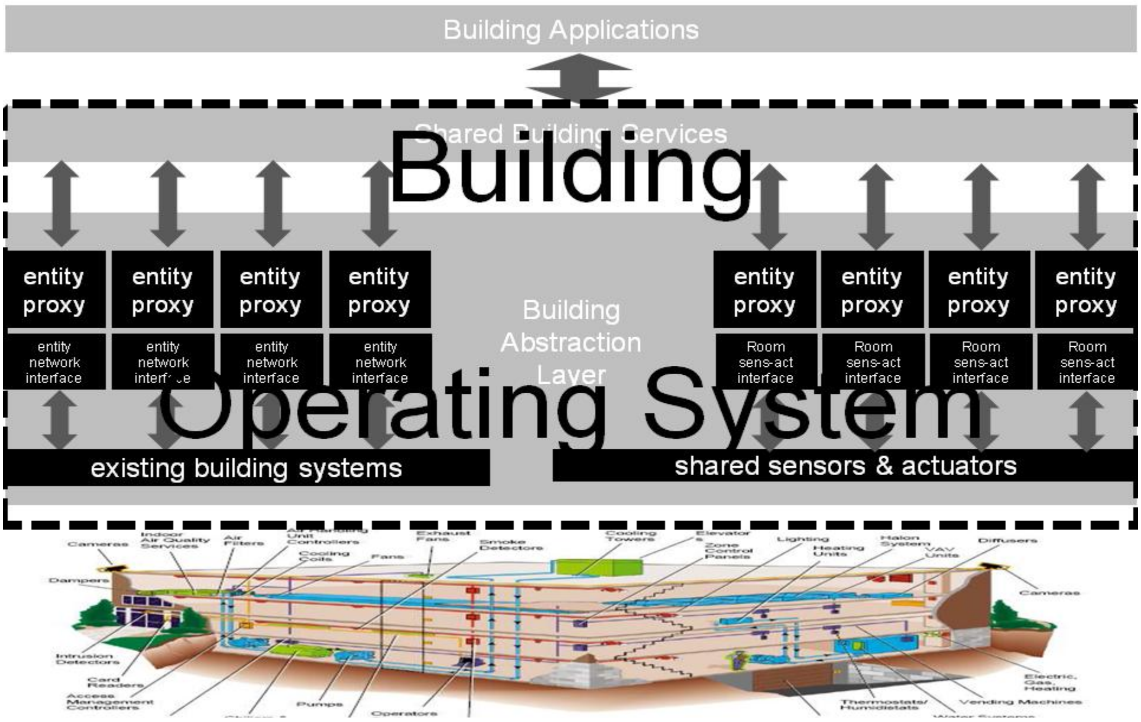

Sensors and actuators are supposed to be shared between all building applications and made available as a pool when they are individually identifiable and addressable (right-hand side of the diagram). Black box legacy systems that do not give access to their individual sensors and actuators will be dealt with through their own interfaces (left-hand side of the diagram).

Building automation applications are not interested in sensors and actuators themselves, but in what is being sensed by the sensors, or acted upon by actuators. The relevant level of abstraction for information pooling should thus be at the level of the physical entities that are being sensed by sensors and acted upon by actuators, which can be pieces of equipment, appliances, people, rooms of a building, or more generally any relevant self-contained subsystems of the building. These entities are generic, intrinsic to the building environment and not tied to any specific building automation application. A set of models and corresponding software components for these entities make up a “Building abstraction layer”, in a way similar to a hardware abstraction layer for a computer platform.

An additional service layer, corresponding to software enablers that span several entities or entity categories, may be provided to applications on top of the building abstraction layer to make up the building operating system. In hte absence of such services, the interfaces that are exposed to applications from this building operating system may correspond directly to the states and associated attributes of relevant physical entities of the building. Taking a room of a building as an example such entity, the state of a room could be whether it is occupied, the type of activity going on, and the corresponding attributes could be its temperature, the number of persons present, etc.. For control purposes, an application can change the state of an entity to another state, if admissible, or change associated attributes. In the examples below, the state of a room could be changed to dark by sending coordinated commands to individual actuators, such as those controlling shades and light fixtures.

The use of this level of information abstraction as a pivotal intermediary layer is in line with the Internet of Things and Context Management enablers provided for the Future Internet platform by the FI-WARE[2] project.

6. Conclusion and perspective

This paper has, on the basis of Building Energy Management Use Cases and ICT requirements identified by the FINSENY project, proposed a view of Smart Buildings as small-scale Future Internet platforms, where a shared building "operating system" can be provided to all building applications, abstracting away the building hardware entities behind the common interface of a generic driver/proxy model.

Prolonging this view, buildings can be seen as semi-autonomous endpoints of the Smart Grid, incorporating into the smart grid an "end to end" principle mirroring the foundational design of the internet itself. As highlighted by the complementarity between local and global optimization use cases, the external interface that smart building energy management systems provide to micrograms and distribution networks should make for proper systems separation of concerns, hiding nested subsystems while providing aggregate monitoring and control entry points.

References

1. *Finseny deliverable D4.1 "Smart Buildings Use Cases Specification" http://www.fi-ppp-finseny.eu/wp-content/uploads/2012/01/D4.1_Smart_Building_Use-Cases-Specification_v1.2.pdf.* 2011.

2. *Overview of ETSI M2M Architecture.* **Pareglio, Barbara.** Sophia Antipolis : s.n., 2011.

[2] www.fi-ware.eu