

GreenMind - An Architecture and Realization for Energy Smart Buildings

Faris Nizamic* and Tuan Anh Nguyen* and Alexander Lazovik and Marco Aiello

Distributed Systems Group, Johann Bernoulli Institute, University of Groningen

Nijenborgh 9, 9747 AG, The Netherlands

Email: {f.nizamic, t.a.nguyen, a.lazovik, m.aiello}@rug.nl

Abstract—Existing buildings are responsible for more than 40% of the world’s total primary energy consumption. Current building management systems fail to reduce unnecessary energy consumption and preserve user comfort at the same time mainly because they are unable to cope with dynamic changes caused by user’s interaction with the environment. To cope with this dynamicity, we propose a software architecture for energy smart buildings that includes a set of concrete software solutions that tackle energy consumption subsystems, i.e., heating/cooling, lighting, workstations, and appliances, in order to save significant amount of energy whilst preserving user comfort. Experimental results carried out in the Bernoulli building, a 12.000 square meter building of the University of Groningen, show that the proposed solutions are able to save up to 56% of electricity used for lighting, at least 20% of electricity used for heating while the savings from controlling workstations as well as other appliances are 33% and 10%, respectively. Totally, our solution is expected to bring up to 28% of saving over total energy consumption in buildings such as the Bernoulli building.

I. INTRODUCTION AND MOTIVATION

Institutions worldwide, such as the European Union (EU), have goals for a 20% cut in Europe’s annual primary energy consumption by 2020. Existing buildings are responsible for more than 40% of the world’s total primary energy consumption [1]. Office buildings are responsible for a significant fraction of the energy consumption and greenhouse gas emissions worldwide [2]. To tackle this issue in Europe, the European Commission proposed several measures to increase efficiency at all stages of the energy chain: from generation to final consumption. The EU’s measures focus on the building sectors where the potential for savings is one of the greatest. However, on 28 June 2013, the European Commission published a report on progress by member States towards Nearly Zero-Energy Buildings (NZEB), which are to become the norm for all new buildings in the EU by the end of 2020. The conclusion of the report is that too little progress has been made by the Member States in their preparations towards NZEBs by 2020 and that Member States have to significantly step up their efforts to implement the requirements regarding NZEBs.

In fact, energy efficiency of buildings is vital because it helps to preserve finite resources, lowers costs for businesses and users, and can be accomplished relatively quickly. In addition, the contribution of low-carbon energy sources is likely to be limited for several decades still. Hence, to move towards a low carbon economy, making “more intelligent” use

of energy in buildings will fundamentally contribute to energy and cost savings.

Building management systems are computer-based control systems installed in buildings that control and monitor the building’s mechanical and electrical equipment. Systems linked to a BMS typically represent 40% of a building’s energy usage; if lighting is included, this number comes close to 70%. BMS systems are thus a critical component for managing energy demand. Improperly configured BMS systems are believed to account for 20% of building energy usage [3], [4]. The general objective of a modern BMS system is to fulfill the occupants’ requirements for comfort while reducing energy consumption during building operations. Heating, ventilation, and air conditioning (HVAC), lighting, hot water, and electricity control are commonly seen as required functions of the BMS system. However, current BMS systems fail to reduce unnecessary energy consumption while preserving user comfort because they are unable to cope with dynamic changes due to users’ interaction with the environment [5]. In addition, heating/cooling and lighting subsystems are usually controlled separately by isolated local feedback loops, thus valuable sensed information, e.g., occupancy patterns, temperature, light level, etc., are not shared and exploited by subsystems. Even more interestingly, computers are even not considered as subsystems of BMS systems leaving as only option for saving energy the use of the default sleep time-out of operating systems (e.g., Windows, Linux, MacOSX). Therefore, to cope with the dynamicity of building’s environment and users’ activity, we propose a software architecture for energy efficient buildings as well as a set of concrete software solutions that tackle separate building subsystems, in an integrated manner.

The developed system manages all building subsystems and optimizes their operations in order to reduce energy consumption while guaranteeing at least the same comfort and productivity for the users. The total integration makes it possible for our system to have as rich inputs as possible. Data is provided not only from all kind of sensors, but also from free, cheap, and already available resources, e.g., computer activity information, weather forecasts, wall material, room size, employee schedules/calendars, room/building reservation systems. The core idea of the solution is to move from isolated local feedback loops to control individual spaces towards an integrated solution that combines improved sensing of the environment and of the user in order to provide a better building experience while saving energy.

The Bernoulli building in the Zernike complex of the University of Groningen serves as our validation and a test-bed.

*The first two authors have contributed equally to this paper, and should therefore be both considered as first authors.

Resources and authorization to work on it came from the Green Mind Award programme of the University of Groningen [6]. The building has a surface of 12.000 square meters accommodating 180 offices, 16 lecture rooms, 8 meeting rooms, and 6 social corners for more than 300 staff members and capacity for more than 500 students. Its annual electric energy consumption is between 1.350.000 kWh and 1.400.000 kWh (for 2011: 1.396.276 kWh and for 2012: 1.359.821 kWh). The experimental results show that our proposed solution is able to save significant amounts of energy. More specifically, our proposed solutions save up to 56% of electricity used for lighting, 20% of electricity used for heating, 33% saving of electricity used for workstations, and 10% of electricity used for other appliances. One should notice that in the overall picture of energy usage, in a typical office building, heating/cooling system accounts for 34% while lighting and office equipment account for 22% and 26% of total energy consumption, respectively [7]. Thus, our solution potentially brings 28% saving in buildings such as the Bernoulli building.

The reminder of the paper is organized as follows. In Section II, we provide a motivating scenario from the user perspective. The general architecture for energy smart buildings is illustrated in Section III. In Section IV, we present our prototype implementation and describe the related pilots. A cost analysis of the system is presented in Section V. In Section VI, we discuss the related work and in Section VII we provide concluding remarks.

II. A WORKING DAY IN AN ENERGY SMART BUILDING

It is the first working day after a long vacation for John, a user of our smart energy building. As John is approaching the building with his bicycle, a sensor at the entrance detects his presence from his office key card. That triggers the heating system to warm up his office to a preferred room temperature and ventilation system to pump in the fresh air. After few minutes, while John is using the stairs to get to his office, his PC bootstraps.

At the moment he approaches his office door, a sensor detects his presence, unlocks the door and turns on the lights. From a room speaker, John hears a pleasant welcoming voice of his virtual secretary: “Good morning, John! It is 9AM and the room is prepared for you. Room temperature is 20 degrees and your PC is on”. While John is hanging his jacket and taking his items out of his bag, he hears more information about his schedule, missed phone calls and important emails to be answered. Besides that, he is also being informed about total energy saving while being on vacation and about new energy saving goals of his department. As he wants to start with his work, he interrupts his virtual secretary by saying: “Thank you, Miranda! That would be all for now”.

At 11AM, John’s colleague knocks his door and enters to give John an update about project developments while he was away. John starts speaking with his colleague and after 3 minutes of inactivity his PC goes to the sleep mode. As he was about to show some vacation photos to his colleague, he gets slightly annoyed by PC sleep action and wakes up the PC by touching the keyboard right away. That directly changes his PC sleep timeout from 3 to 5 minutes.

Now, it is 1PM and John leaves his office for a lunch. The room sensors detect his absence and turn off lights, heating and ventilation, and put his PC to sleep mode. Now, for one hour John does not consume any energy in his office. As John arrives first to the building restaurant, lights go on only in the area where he sits, while other lights in neighbouring areas are dimmed. While he is enjoying lunch with his colleagues, he notices that lights above them also get dimmed as the natural light increases. As soon as he is finished with his lunch, he passes an occupancy detection sensor at the staircase and the room preparation actions are triggered again. Only this time after entering his office, the lights do not turn on. The reason for that is a light sensor detected enough natural light coming from outside. Also, heating did not start 10 minutes before entering the office, but only 3 minutes before. The reason for that is because both outside and inside temperature sensors have detected a significant increase in temperature, caused by the very sunny weather. This results in energy saving and John is informed about it by his virtual secretary. Motivated by that, to contribute to the saving goal of his department, John decides to cool himself down a bit and asks his virtual secretary: “Miranda, please decrease heating of 20%. Thanks”.

At 5.45PM, John finishes with his work and leaves his office in a rush to get on time to a dining place. Sensors detect his absence and turn off all energy consuming devices immediately. John’s office again becomes energy neutral until the next working day.

III. THE GREENMIND ARCHITECTURE

Towards the realization of the just introduced scenario, we propose an architecture that takes, as its inputs: (1) user activities, such as working with PC, being present, absence, (2) appliance statuses, e.g., PCs are idle for a certain amount of time, and (3) environment information (natural light intensity, outdoor temperature, etc.). The system accordingly adjusts the environment to preserve user comfort and to save energy. The proposed architecture goes from the physical level of consumption measurement, live environment monitoring, up to Hierarchical Task Network (HTN) based controlling. Furthermore, the system also provides displays that bring consumption information at the building level, but also at the personal level.

A. Design Principles and Overall Architecture

The governing of the system consists of multiple sub-components. Components work together to reach the common energy saving goal, sharing useful information with the other components so that both the raw sensor data and other processed information is accessible to each component. To make the management of these components easier, the components should be loosely coupled and the communication should be asynchronous.

It is envisioned that over time new sensors and actuators can be added to the system, as well as new components that provide better or completely new functionality, so that the system keeps improving. Through the loose coupling of the individual components the system scales better, it adapts more easily to new components, and is easier to manage than a conventional BMS. By using a dynamic combination of small hardware, that integrates with the existing building design, and

dedicated software the system has the potential to save energy at a relatively low cost without the loss of comfort and even the potential to increase comfort.

The overall architecture of the system, named *GreenMind*, is shown in Figure 1. At the bottom lays the *Physical layer* that contains all devices connected to *Sensor and Actuator Gateways* (SAGW). The Physical layer also includes *Sleepy*, a software agent that runs on clients, i.e., workstations, working both as sensor - for monitoring and actuator - for controlling workstations. The data processed by SAGW is very important for the *Context* component to reason and provides a complete and consistent view of the environment. This essential view is stored in a *Central Database* together with consumption data from the *Consumption Measurement* component. Also, the crucial view about the environment is provided to the *Controller* that makes decisions how to control electricity consumers. The control decisions made by the controller are sent to the *Orchestrator* that translates them into proper commands distributed to either SAGW that takes care of actuation or the *Sleep Management Solution - SMS* component that takes care of control of workstations through interacting with the *Sleepy* software agent. In upper right part of the architecture one can find the *Dashboard/Mobile App* component that is responsible for delivering, in a very meaningful way, the consumption information to each and every occupant of the building.

B. Physical Layer and Gateways

At the Physical Layer, sensor networks provide the basic infrastructure for gathering raw contextual data from the environment. For example, light and temperature sensors gather information about ambient context while passive-infrared (PIR) and pressure sensors or electricity measuring plugs provide additional insights.

1) *Gateway*: The environment is being interpreted through information gathered by sensors. The more sensors are deployed, the more precise information can be deduced. The diversity of sensors implies that the system should be able to address the heterogeneity of sensor networks. Thus, the Gateway is capable to read information from different networks and join data read from them into only one standardized message with a specific format, realizing the first step of data processing. In addition, the Gateway encapsulates the complexity of sensor networks and makes components independent from each other.

2) *Interface to Upper Components*: We use a message broker to reduce coupling between low-level components (i.e., sensor networks and Gateways) and the high-level ones (i.e., Context), which makes it possible for the components to run independently in a distributed manner.

C. Context Component

The Context takes as its input the sensor data processed by the Gateway. The Context component has two subcomponents: (1) *ambient context* and (2) *activity recognition*. The ambient context takes care of other ambient information, such as natural light intensity. Furthermore, we employ context consistency diagrams (CCD) [8] as a key component for fault correction. A CCD is a data structure that provides a mechanism for probabilistic reasoning about the current situation and determines the most probable current situation in

the presence of inconsistencies, conflicts, and ambiguities in sensor readings. With reliable context information provided by CCD, for user activity recognition, we apply an ontology-based activity recognition approach to derive high-level activities from data from simple sensors. Our ontology-based activity recognition solution is represented in [9]. For more detail about the inside implementation of our context component, one should refer to [10].

D. Fine-Grain Energy Consumption Measurement

We measure consumption at the device level, i.e. electric energy consumption by an occupant. Besides occupants' devices (e.g., PC, mobile phone charger, water cooker), we also measure the electric energy consumption of light. That way we know total consumption per user. This enable keeping of the user consumption history, as well to give user a feedback on the consumption.

E. Database

The database is necessary to store and retrieve the data (e.g., the sensor measurements) for the components of the system. The data can be stored in relational databases using tables, in a graph database, or both. The primary difference is that in a graph database the relationships are stored at an individual record level, whereas in a relational database the structure is defined at a higher level (i.e., table definitions). One of the main advantages is that graph databases provides index-free adjacency (i.e., graph traversals can be performed with no index lookups) that can lead to a much better performance. Each sensor measurement contains relationships to other data such as the floor, building, time and date, and room or user. As a large amount (e.g., 1,000,000) of sensor measurements with a traditional database caused severe performance issues (e.g., retrieve the total energy consumption for a certain hour on a specific date), thus a decision is made to use a combination of NoSQL database and graph database. NoSQL database delivers faster performance for systems involving time series or big data (such as data coming from thousands of sensors), while graph database has advantages for maintenance of building structure and relationship between nodes (e.g., users, rooms, devices).

F. HTN-based Controller

The components belonging to the composition level use all the available data and functionality from the lower layers to take decisions about the desired state of devices. This is also the layer where the system management components are located. The controller is the component responsible for providing the system with the ability to automate the sensor and actuator behaviour, which influences the energy consumption. We propose an approach that manipulates context information, and combines an ontology-based technique to recognize occupant activities with Artificial Intelligence (AI) planning to automatically produce office adaptations at run-time. The proposed approach is initially described in [11]. Given the context information and the recognized activities, Hierarchical Task Network (HTN) planning is adopted to compose a plan, that is, a sequence of device operations. Then, plans are executed by using a network of actuators. The HTN planner takes care of controlling lighting systems and HVAC, but it also invokes the Computer Sleep Management (CSM)

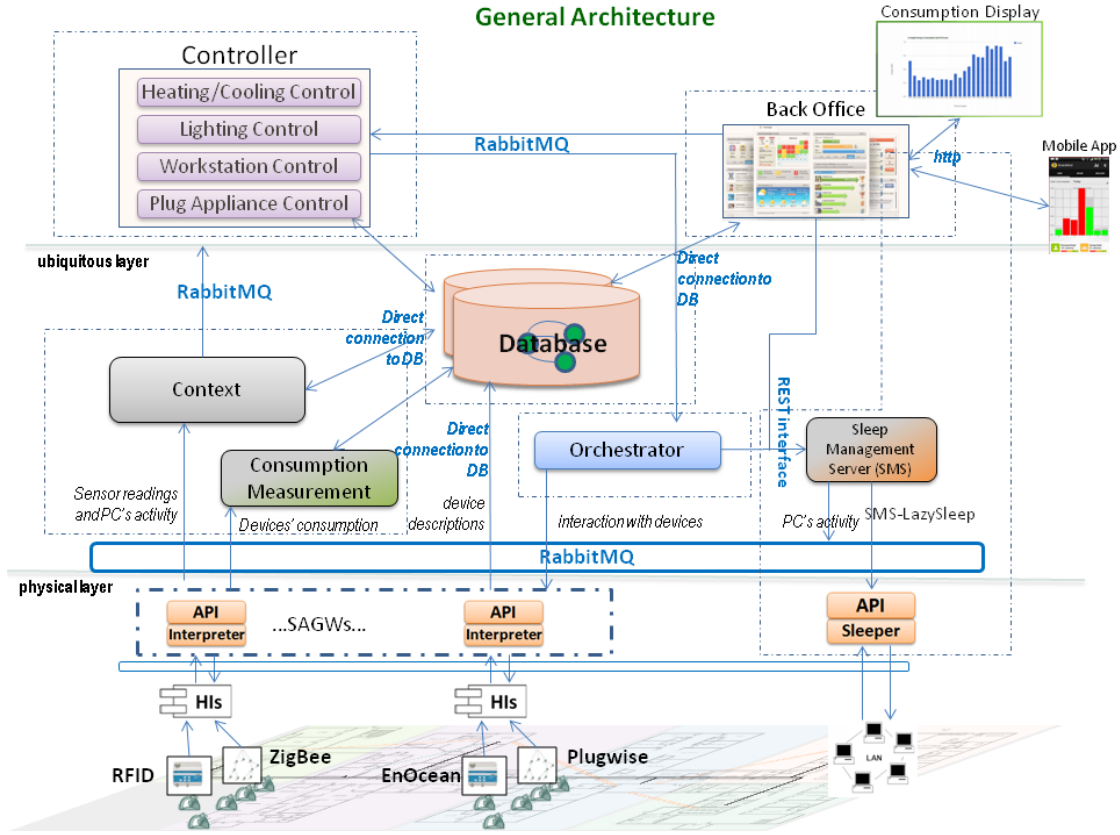


Fig. 1. GreenMind system architecture

component for controlling workstations in order to reduce energy consumption while not hindering an user by providing personalized workstation sleep timeout.

G. Computer Sleep Management

The goal of the Computer Sleep Management (CSM) solution is to minimize the energy usage of workstations. The way CSM achieves this goal is by taking control of the process for putting a workstation into sleep mode when no activity is detected. The CSM also provides the administrators with important information with regards to the activity history of a workstation and whether a workstation should be sleeping or not. Besides this, using the CSM Dashboard GUI we can see the activity of a user and put workstations to sleep, wake them up, turn them on or off; or simply view the percentage of time when a PC was off, working or in sleep mode.

The CSM consists of two components: (1) Sleepy Client, a client monitoring power management events and workstation activity while listening to and executing commands from the server, and (2) Sleep Management Server, a server managing data received from all clients and listening to requests from other components, forwarding these requests to the clients. More technical details on CSM solution are described in [12].

H. Display

To raise awareness and decrease energy consumption as much as possible it is important to bring users into the loop. The GreenMind system is able to provide personal energy

consumption information to each user in a very meaningful way, thus motivating users to save energy consumed by their appliance. The graphs are presented with varying level of granularity, more specifically: in daily, weekly and monthly view. GreenMind provides mobile application as well as web-based interface with which users are able to keep track of consumed energy, to set saving goal for a certain period (e.g., week, month), to participate in saving campaigns managed by building managers, etc. We also show more meaningful information such as a price of consumed energy as well as an amount of emitted CO₂.

I. Orchestrator

To keep the controller as simple as possible, the execution of the plans made by the controller is left to the orchestrator. The orchestrator component acts a buffer between the physical layer and the controller, and it takes care of sending the commands at the correct time, as well as the order of execution. Each plan can be started and stopped at any time, or paused and resumed when necessary. To limit the chance of the orchestrator component becoming a bottleneck, the orchestrator is multi threaded and implemented in such a way that multiple instances can execute simultaneously.

The orchestrator communicates with the SMS and SAGW, depending on the instructions received from the controller. The SMS handles the requests from the orchestrator to change the state of workstations and the SAGW handles the state change requests for the other devices. The current state change

requests are straightforward and consist out of a unique identifier for the device (e.g., a MAC address, IP address) and the desired status of a variable. The status can be on, off, sleep, or hibernate, in case of a workstation.

J. System Communication

System communication is based on message queues. These are used by the components of the system to communicate with each other, with or without knowing the physical location of other components. Message queues can significantly simplify the implementation of the separate components and also improve performance, scalability and reliability. Another advantage of using a message queue is that the sender and receiver do not need to communicate with the message queue at the same time as the messages are stored onto the queue until the recipient receives the message. As a result, the communication between different components is asynchronous.

K. System Operation

The typical operation cycle is as follows: 1) the SAGW and SMS send messages to the context based on the changed values from their sensors and actuators, 2) the context component processes these messages to create new variables and sends the variables values to the controller component, 3) the controller executes HTN planning based on the received values, 4) messages that require action are then sent from the controller to the orchestrator, 5) the orchestrator then sends commands to either the SMS or SAGW, depending on the received message, and 6) the SMS or SAGW then execute the requested action (e.g., turn off a device or put a workstation to the sleep mode), and then the cycle repeats. The components do not wait until a full cycle is completed (e.g., a change in the sensor value can occur at any moment).

IV. PROTOTYPE IMPLEMENTATION AND PILOT PROJECT

A. The GreenMind Prototype Implementation

We realized a prototype implementation and deployed it in the Bernoulli building. In the prototype, we implement a gateway that can handle a wireless network of sensors and a wireless network of electricity measuring plugs, providing enough ability to monitor essential environment information. In addition, we consider the Sleepy agent as another type of sensor and actuator that can monitor the activity of computers and can control their sleep mode. The first version of Back-Office, Display, and Mobile application are released, together with a complete implementation of HTN-based controller, ontology-based context component.

1) *Physical Layer and Gateway*: In the current implementation, we use IEEE 802.15.4 compliant wireless TelosB-based sensors produced by Advantix Systems [13]. The on-board PIR is used together with the light sensor. The motes are programmed in nesC and run on the TinyOS 2.1.1 embedded operating systems. For electricity measuring plugs, we use Plugwise [14] products consisting of plug-in adapters that fit between a device and the power socket. The adapters can turn the plugged device on and off, and can, at the same time, measure the electric energy consumption of the device attached. The plugs form a wireless ZigBee mesh network around a coordinator. The network communicates with the base

station through a link provided by a USB stick device. We implement the Gateway as a process running in the background that reports power state of controlled devices. The gateway is written in Java.

2) *Interface to Upper Components*: The RabbitMQ [15] messaging framework is chosen as it is a complete and highly reliable enterprise messaging system based on the emerging AMQP [16] standard and runs on all major operating systems.

3) *Consumption Measurement component*: The Consumption Measurement (CM) component gathers data from the sensor and actuator gateway (SAGW) and stores it in the database Neo4j [17]. The CM is written in Java. It uses the history buffer to limit the network traffic and processing load. Requesting all history logs until the last log that was stored helps in case of CM failure. The Plugwise devices continue to monitor the electric energy consumption and when the CM restarts it simply gathers all the missing data. The same also applies when the network of the Plugwise devices is temporarily unreachable. When the network connection is re-established the consumption component will retrieve all data until the last stored history log.

4) *Context*: The ontologies are developed using Protégé [18], a graphical tool for ontology development that simplifies design and testing. Ontological reasoning is performed using the HermiT [19] inference engine, and its application programming interfaces (APIs) for the Java programming language. The recognition algorithm is developed in Java and implemented as an on-line recognition system. More details are presented in [9].

5) *Databases*: In our system, we use two databases, Neo4j and Cassandra. Neo4j is an open source database with features from both document and graph database systems. Neo4j excels in scalability, availability, performance, and price. After careful consideration, Neo4j is picked as the main storage facilitator for the building spatial and structural information of the GreenMind system. Cassandra [20] database is a NoSQL database that delivers fast performance for systems involving time series or big data. Cassandra is used to store time series data coming from sensor readings periodically (e.g., each second one data point).

6) *Computer Sleep Management*: The Sleepy client is written in Scala, the other components at the sever side are written in Java. At the moment, the Computer Sleep Management supports Windows XP and Windows 7. Current development focus is on features for Windows 7 as that is operating system that is most widely used at the University of Groningen and also supported by its IT services. Other operating systems, such as Linux and Mac OS, are left as future work. The Sleepy runs on workstations and is responsible for monitoring the activities on them. Currently, an activity is defined as keyboard and mouse movement. Sleepy is also responsible for changing the sleep timeout (time before a workstation enters sleep mode after being idle/inactive for a specified amount of time), as well as executing "turn off" and "go to sleep" commands received from the Sleep Management Server. The Sleep Management Server is responsible for collecting data from Sleepy clients and publishing this data on RabbitMQ. Furthermore it is also responsible for forwarding commands for Sleepy clients to the correct Sleepy client by mapping MAC addresses to IP

addresses. For more details about the implementation of the Computer Sleep Management we refer to [12].

7) *HTN-based Controller*: All components of the Planning System are implemented in the Scala programming language. The problem converter translates the context information and activities described in JSON syntax into Hierarchical Planning Definition Language (HPDL) syntax [21]. The domain modeller enables creating models. Consequently, the planner's input consists of HPDL domain and problem descriptions. Planning is offered as a service by implementing its functionalities as REST resources. Upon receiving a request with appropriate arguments, the planner searches for a solution, structures the resulting plan in JSON, XML or plain-text format, and returns it to the interested party.

8) *System Communication*: All the components in the GreenMind system communicate with each other through the use of JSON objects. JSON was picked because of the good readability, simple syntax and ease of use. JSON is also less verbose than alternatives like XML, which decreases the size of the message payloads. Developers of components that communicate with each other dictate the content of the messages that are communicated with each other. Depending on the data that the component requires, the body of the JSON object may change. For the sake of traceability with asynchronous communication and readability and consistency in generic messages, a basic template was specified to which all the messages of components should adhere.

B. Pilot Project Description and Setup

We have developed and tested, including a number of energy saving solutions, including: 1) *Consumption measurement*, 2) *Consumption display*, 3) *Computer sleep solution*, 4) *Lighting control*, 5) *Heating/cooling control*, and 6) *Plug appliances control*. In this section, we give more details on the deployment of each solution within the pilot. Then, we illustrate the setup of the experiments that we use to evaluate our proposed solution and the results, especially in terms of energy saving.

1) *Consumption Measurement Lab*: To understand individual energy consumption and its patterns, we deployed electricity measurement devices in 12 areas on the 5th floor of the Bernoulli building. Layout of the consumption measurement lab is illustrated in Figure 2. The areas that are measured by the system include eight offices, two hallways, one kitchen area, and one meeting area. Over a course of six weeks, from mid-May to late July, 2013 the office environment on the fifth floor of the Bernoulli building has been monitored 24 hours a day and 7 days a week.

The information from Consumption Measurement was used to develop the BackOffice component which serves to model building spacial data and to assign particular measurement devices to users, users to rooms, etc.

2) *Consumption Display*: To understand the pattern of electric energy consumption at the level of the building as a whole, we developed a consumption display and deployed it at the Computer Science social corner on the 5th floor of the Bernoulli building, Figure 3. The information shown on the consumption display consists of daily and weekly view of



Fig. 2. Consumption Measurement Lab




Fig. 3. The Consumption Display at the 5th floor of the Bernoulli building

consumption, together with moving average marks showing average value for past four weeks. In order to show this information on the display, we had to use external hardware component to convert analogue light signal from the main electricity meter to a digital signal and to store it to a database before being able to represent it graphically.

3) *Computer Sleep solution*: In order to understand the patterns of workstation usage, sleep timeout preferences and part of electric energy consumption coming from workstations, we have deployed Computer Sleep solution to 10 workstations at the 5th floor of the Bernoulli building. Per monitored workstation, we show the activity (working or idle) and status (off, on, sleep). Besides that, we can control the workstations using four actions: 1) send to sleep, 2) wake up, 3) turn off, and 4) turn on. The activity and status of workstations are shown on a dashboard, see Figure 4. Also, control options can be used via the dashboard. Historical data is also stored in our database and it can serve to determine users' consumption using workstation usage time and average workstation electric energy consumption values. Through database, we can generate many meaningful reports which may serve to motivate further energy preservation. Currently, we are working on the deployment of a sleep solution to the about 300 PCs present in the building.

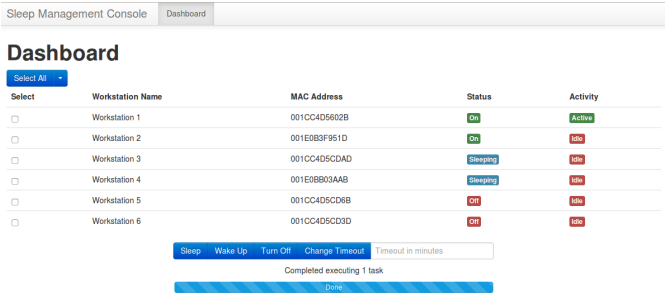


Fig. 4. Computer Sleep solution - Control Dashboard

4) *Lighting control*: The Lighting control system is deployed in the Restaurant on the ground floor of the Bernoulli building. There, we attached sensor/actuators to each individual lamp (30 in total), see Figure 5. That way, we can both measure the consumption as well as control each individual lamp or a group of lamps. Besides that, we also deployed motion and light sensors which give the input to the controller component to decide when lights do not necessarily need to be used (i.e. when no one is present in a particular area or when natural light provides satisfying high level of light). Using this system, we reduce electric energy consumption by using lights only when they are absolutely necessary.



Fig. 5. Placement schema of the 30 lamps and plugwise devices in the restaurant at the ground floor of the Bernoulli building

5) *Heating/cooling control*: The goal of heating/cooling control solution is to reduce energy consumption of heating/cooling system by reducing heating/cooling time within room preparation to minimal required. The experimental results show that very high savings can be achieved. The heating/cooling control solution is tested in separate premises outside the University of Groningen. One of the reasons for that is the need to test the solution in an environment that is highly controllable before applying it to the production environment. Experiments and results of our proposed heating/cooling control is represented in Section IV-C. Control of the heating/cooling system was done under the assumption that heating and ventilation can be separately controlled at the room level. The setup of our experiment is shown in Figure 6. The experiment was carried in one week, from Jan 27 to Feb 2, 2014. In this experiment, we used: a computer to run the prototype software, a heater with capacity of 2KW (for safety reason, the heater was set to perform at 25% of its full capacity), a Meteo sensor that can measure humidity, pressure and temperature of the environment, a Plugwise device to control the heater [22].



Fig. 6. Heating/cooling experimental setup

C. Energy Saving Evaluation

We ran intensive experiments at our pilot project areas to evaluate our solution in terms of accuracy, scalability, and, most importantly, in terms of energy saving. In the current treatment, we only consider the last item as the accuracy and scalability have been covered in previous publications [9], [10], [11]. Incidentally, we remark that the system itself consumes energy to operate. In the current implementation, we use TelosB-based sensors produced by Advantix Systems. The TelosB sensors consume insignificant amount of energy (5.4 mW in active and 15 μ W in sleep mode). Currently, we are working to replace them with energy harvesting wireless sensors, such as EnOcean solar or motion energy harvesting ones. Thus, additional electric energy consumption for our sensors is close to zero. For computer sleep management, no additional sensors or devices are required. For plug appliances control, we use Plugwise devices which consume a maximum power of 1.1 W; this is insignificant with respect to the overall consumption of the controlled devices. To collect sensor readings and to send actual commands to the actuators, we use two energy-efficient thin clients which consume a negligible 15 W per piece. Additionally, our integrated system can run on a shared sever that is certainly always on to provide other services to the users, so it could add insignificant consumption to the already active computers. Lastly, in a real operational environment, the system would manage many more devices; thus, its energy consumption would be amortized over larger savings. For these reasons, we have not included these energy consumptions in the current evaluation.

1) *Accuracy and Scalability Evaluation*: Regarding the accuracy and scalability of the solution, each and every proposed approach and component was tested in our previous research, in which experiments were conducted at our living lab at the University of Groningen. Experimental results are discussed in detail in our previous publications. In particular, in [9], the experiments conducted show that in such experimental setup, our ontology-based activity recognition solution is able to recognize six typical office activities (working at a desk with or without a PC, having a meeting, having a coffee break, and presence/absence) in three office activity areas (two working rooms and a coffee corner) for two persons with average accuracy over 88%. In [10] we evaluated the performance of CCD when applied to correct faults in sensor readings. Results show that CCD helps to significantly improve the accuracy of the context component, thus producing reliable context

information from unreliable sensor data. The performance of the HNT-based controller is evaluated extensively in [11].

2) *Energy Evaluation:* With respect to energy saving evaluation, in [23] we discuss the potential saving brought from plug load control with experiments conducted in two weeks. The results show that our solution saved 10% of electricity consumed by a set of appliances (a fridge, a laptop, a printer, a projector, a microwave, and a water boiler). Meanwhile, saving achieved from controlling lighting and workstations was initially evaluated in [11], where we run experiments within two private offices of two users and determined how much electricity saved by controlling their two workstations and ceiling lamps within their offices. The overall solution shows intriguing potential for energy saving in the order of 70%, given mostly sunny days and a provisional set of devices for experimentation [11].

Separate experiments conducted for each individual solution show that by applying our solutions it is possible to achieve significant savings from all four main subsystems (heating/cooling, lighting, workstations, and plug loads) within office buildings.

3) *Energy Saving Potential:* In the following, we present analysis on energy saving that is achieved by using our GreenMind system. The current estimated annual values are calculated based upon the average energy consumption of the actual measured devices within and outside the office hours during six weeks, whereas the after estimated annual values are estimated only based on the possibilities of the GreenMind system to limit the energy consumption.

a) *Potential Saving from Controlling Lighting System:* Currently, the most energy-efficient lighting solution is LED lighting. In addition, most of the lighting solutions use Passive Infra-red (PIR) and light intensity (LUX) sensors to control and dim lamps based on movement and natural light. However, such systems work in a closed loop, that means the inputs of the system come from the PIR and LUX sensors only. Furthermore, the time out is static and usually set manually during the installation/calibration phase. Our solution brings the dynamicity required for the system in order to save more electricity while preserving user comfort. With our solution, lighting would be adjusted according to specific user activity (e.g., working with PC, reading, having a meeting, etc.), thus user comfort and productivity would be enhanced.

Particularly in our realization within Bernoulli building, an existing lamp fixture is highly energy inefficient and it consumes around 75 W when turned on. There are two or three fixtures inside each office. This means that the lighting of a single office consumes from 150 W to 225 W hour for around 8 hours. This equals to approximately 1.2 kWh to 1.8 kWh per day, that is simply being wasted on a sunny day. According to a weather condition report from the Royal Netherlands Meteorological Institute [24], there is an average of 4.5 hours (56% of the eight hours of work time per office) of sunlight per day and 228 workdays per year, the annual savings could amount to 153.9 kWh to 227.81 kWh per office annually. Specifically, in an extreme case we evaluated in [11], the overall solution shows intriguing potential for energy saving in the order of 70%, given mostly sunny days and a provisional set of devices for experimentation.

b) *Potential Saving from Controlling Heating/cooling System:* In principle, our solution is first of all based on the ability to control individual heating systems at the room level. We use wireless single room control module called *en:key* from Kieback&Peter [25], which currently is one of the most-advanced technologies for heating/cooling control using machine learning. Our solution improves it by optimizing the heating of rooms based on the fusion of information regarding users presence, external weather conditions, weather forecast, and university/user calendars. The goal is to control the room heating automatically and heat only when users are present, thus, optimizing the temperature given the weather exposition and the actual number of occupants present in the room.

Despite of the fact that we could not test this component within Bernoulli building, we collected interesting information from the running of the component elsewhere. In fact, the experimental results show that the amount of time that the heater needed to warm up the room differed every day. The duration for preheating changed based on the initial temperature of the room, which is affected by the outside temperature. In the worst case, the heater needed to be turned on up to 80 minutes before the scheduled time in order to warm the storage house to the aimed 10 degrees Celsius. Yet, for some days the heater was just turned on approximately 18 minutes before the scheduled since it was a warm day. The difference between two above days is almost 60 minutes in using the heater, which results in energy saving.

In either case, the experimental results are very satisfying as the clear relation between outside weather and necessary heating start time was noted. Additionally, by applying machine learning algorithms on top of this solution, we could more precisely determine latest heating start time. Finally, depending from the schedule of the room, our experiments showed that electricity savings can be 20% or even higher on heating, depending on the room schedule [22].

c) *Potential Saving on Controlling Workstations:* A standard university workstation consumes 120 W when both computer and LCD screen are on. The consumption slightly decreases around 13:00 when staff members typically go for lunch and manually turn off the monitor and/or computer for the duration of the lunch. Note that the workstation still consumes some energy outside office hours.

The estimation of the energy savings potential is of at least 33% per regular workstation. This comes from a survey conducted with 15 staff members who are occupying our living on the 5th floor. The survey shows that on average, a staff spends behind the computer around five and a half out of the eight working hours, meaning that for about two and a half hours the computer can be sent to sleep mode in order to save energy. More detailed discussions are presented in [10].

d) *Potential Saving from Controlling Other Appliances:* Based on the measurements collected from a microwave, a coffee machine, a water boiler, and a fridge, we spot the potential saving by controlling other appliances within office environment. Potential saving achieved is around 10%. One can also find detail discussions on the potential saving from controlling other appliances in [23].

For example, the fridge is cooled on intervals and can not be turned off completely because the drinks and food in the

fridge need to be cooled to stay fresh. The experimental results in [23] show that the fridge is capable of cooling sufficiently when only turned on for 15 minutes every hour. During these 15 minutes the fridge is cooling at the maximum energy level. Taking into account that the rated electric energy consumption of the fridge is 70 Wh, the consumption of the fridge during those 15 minutes per hour should be around 17 Wh at most.

As a result, there is a difference of 65 Wh between a regular weekday and a day in the weekend. This would mean a saving of 325 Wh (5 days * 0.065 kWh) per working week and a yearly saving of 14.82 kWh (228 working days * 65 Wh) per fridge, which amount to around a potential energy saving of 10%.

V. FINANCIAL CONSIDERATIONS

A. Installation Costs

From a financial point of view, to realize and deploy the proposed solution for a generic building, we need to equip all rooms (e.g., private offices, lecture rooms, social areas, cafeteria) with sensors and actuators. Also, a server is needed to run the software of the architecture. For ease of cost calculation, we estimate the system cost based on the number of rooms within a building. In side each room, we need to have 1) a set of sensors gathering information about movement, light intensity, temperature, consumption and 2) a set of actuators controlling light, heating/cooling, workstations, and appliances. We estimate the average cost per room to be 200 euros, if we want to have a full set of sensors/actuators, which is not always the case, for example the cafeteria does not contain workstations. Using the Bernoulli building case study, we note that it has 200 rooms. Therefore, the cost is approximately 40.000 euros.

B. Break even point

Considering a potential saving in the energy consumption for the Bernoulli building at about 7%. Given that the annual electric energy consumption of the building is approximately 1.400.000 kWh, then, the annual saving is 98.000 kWh. This saving is equivalent to 11.500 euros per year, at current energy prices. Thus, the break even point for the Bernoulli building comes in less than four years. Taking into account that the potential saving can be even higher (up to 28%), such time can be significantly reduced. If we consider that the price for hardware/sensors are getting cheaper while the energy price is increasing, in the future, the break even point can occur even more rapidly.

VI. RELATED WORK

Several studies exist on energy saving potential in office buildings. Lighting has been the focus of many authors [26], [27] as it is an area with significant energy improvement potentials. Besides lighting, office equipment (e.g., computer and monitor) influences the energy usage to a great extent. Kawamoto et al. [28] and Webber et al. [29] show that a large percentage of computers and monitors are not turned off but in standby mode during non-office hours (40% in Japan and 75% in US). According to the authors of [29], computers and monitors in offices are found to be in standby mode for no

less than 21 hours a day. HVAC systems can also contribute to the possible energy and cost savings according to several studies [30], [31]. Our solutions focus on all subsystems, i.e., lighting and office equipment for energy saving potential as well as the HVAC.

In [5], we describe in a survey about energy intelligent systems that the three main energy consuming subsystems in buildings (i.e., HVAC, lighting and office equipment) also draw the attention of numerous studies. From the researched studies, one notices that some focus on one subsystem only, while others try to save energy for two or even three subsystems. The majority (22 out of 32) of the reviewed studies focus only on one subsystem, our solution focuses on more subsystems. Six studies focus on HVAC and lighting, three focus on HVAC, lighting and power plugs, and two others focus on lighting and power plugs only. The iSpace [32] and the Intelligent Buildings by Davidson et al. [33] are the two projects that only focus on lighting and power plugs, similar to the focus our research. The work presented in [32] differs with our research because, although it can be used for energy savings, it is focused on autonomous environments and on student bedrooms, rather than offices. In addition, the energy consumption is not monitored.

VII. CONCLUSIONS

The GreenMind system is capable of reaching the goal of saving 7% of energy, and it has the potential to save up to 28% in buildings such as the Bernoulli building. All subsystems, i.e., lighting system, heating/cooling system, workstations, and other appliances can be monitored and controlled. Experimental results confirm that significant amount of energy spent can be saved. Specifically, 56% is the electricity saving from lighting control, the saving from heating/cooling control could reach at least 20%, while the same figures for workstation control and plug appliances control are 33% and 10%, respectively. Furthermore, the GreenMind system is a framework that allows the University of Groningen to expand its research into sustainability in a new direction, such as user awareness, and provides the researchers with real and up-to-date data. The loose design of the system components means that it can be enhanced over time, exceeding the goals of this initial design, and thus provides more possibilities than the existing built-in BMS.

ACKNOWLEDGEMENTS

The authors thank Brian Setz, Ruurtjan Pul, Mattijs Meiboom, Sijmon Heitmeijer, Marcel Jillings, Rosario Contarino, Diederik Jan Lemkes for contributing to the implementation of the system. This work is supported by the EU FP7 Project GreenerBuildings, Contract No. 258888, the Dutch National Research Council project Energy Smart Offices, Contract No. 647.000.004, and by the University of Groningen Green Mind Award. Faris Nizamic is supported by the European Commission within the Erasmus Mundus Action 2 programme (JoinEU-SEE). Tuan Anh Nguyen is supported by the Vietnam International Education Development Program.

REFERENCES

- [1] J. C. Howe, "Overview of green buildings," *National Wetlands Newsletter*, vol. 33, no. 1, 2010.
- [2] J. Borggaard, J. A. Burns, A. Surana, and L. Zietsman, "Control, estimation and optimization of energy efficient buildings," in *Proceedings of the 2009 Conference on American Control Conference*, ser. ACC'09. Piscataway, NJ, USA: IEEE Press, 2009, pp. 837–841. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1702315.1702453>
- [3] M. R. Brambley, *Advanced sensors and controls for building applications: Market assessment and potential R&D pathways*. Citeseer, 2005.
- [4] K. W. Roth, D. Westphalen, J. Dieckmann, S. D. Hamilton, and W. Goetzler, *Energy Consumption Characteristics of Commercial Building HVAC Systems: Volume III, Energy Savings Potential*. National Technical Information Service (NTIS), U.S. Department of Commerce, 2002.
- [5] T. A. Nguyen and M. Aiello, "Energy Intelligent Buildings based on User Activity: a Survey," *Ener. and Build.*, vol. 56, pp. 244–257, 2013.
- [6] "Green mind award website," <http://www.rug.nl/about-us/who-are-we/sustainability/green-mind-award>, 2013.
- [7] "U.S. Department of Energy, 2010 Buildings Energy Data Book," 2011.
- [8] V. Degeler and A. Lazovik, "Reduced context consistency diagrams for resolving inconsistent data," *ICST Tran. on Ubiquitous Env.*, vol. 12, no. 10-12, 2012.
- [9] T. A. Nguyen, A. Raspitzu, and M. Aiello, "Ontology-based office activity recognition with applications for energy savings," *Journal of Ambient Intelligence and Humanized Computing*, pp. 1–15, 2013. [Online]. Available: <http://dx.doi.org/10.1007/s12652-013-0206-7>
- [10] T. A. Nguyen, V. Degeler, R. Contarino, A. Lazovik, D. Bucur, and M. Aiello, "Towards context consistency in a rule-based activity recognition architecture," in *10th Int'l. Conf. on Ubiquitous Intelligence and Computing/Autonomic and Trusted Computing (UIC/ATC)*, Dec 2013, pp. 625–630.
- [11] I. Georgievski, T. A. Nguyen, and M. Aiello, "Combining activity recognition and ai planning for energy-saving offices," in *10th Int'l. Conf. on Ubiquitous Intelligence and Computing/Autonomic and Trusted Computing (UIC/ATC)*, Dec 2013, pp. 238–245.
- [12] B. Setz and R. Pul, "Lazy sleep," University of Groningen, Internship Report, jul 2013. [Online]. Available: http://www.cs.rug.nl/~faris/students/Brian_Ruurtjan_report.pdf
- [13] "Advantics Sys." <http://www.advanticsys.com/>, 2013.
- [14] "Plugwise," <http://www.plugwise.com/>, 2013.
- [15] A. Videla and J. J. Williams, *RabbitMQ in Action*. Manning, 2012.
- [16] "AMQP Messaging System," <http://www.amqp.org/>, 2013.
- [17] "The neo4j database," <http://www.neo4j.org/>, 2014.
- [18] "Protégé," <http://protege.stanford.edu/>, 2013.
- [19] "Hermit Reasoner," <http://hermit-reasoner.com/>, 2013.
- [20] "The cassandra database," <http://cassandra.apache.org/>, 2014.
- [21] I. Georgievski, "HPDL: Hierarchical Planning Definition Language," Uni. of Groningen, JBI Preprint 2013-12-3, 2013.
- [22] L. D. Tuan, "Efficiency energy usage by dynamic scheduling the hvac system," University of Groningen, Internship Report, feb 2014. [Online]. Available: http://www.cs.rug.nl/~faris/students/TuanLuu_report.pdf
- [23] I. Georgievski, V. Degeler, G. A. Pagani, T. A. Nguyen, A. Lazovik, and M. Aiello, "Optimizing energy costs for offices connected to the smart grid," *IEEE Transactions on Smart Grid*, vol. 3, pp. 2273–2285, 2012. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6377248>
- [24] "Royal netherlands meteorological institute," <http://www.knmi.nl/klimatologie/>, 2013.
- [25] "The en:key solution from kieback&peter," <http://www.enkey.de/>, 2013.
- [26] M.-C. Dubois and Å. Blomsterberg, "Energy saving potential and strategies for electric lighting in future North European, low energy office buildings: A literature review," *Energy and Buildings*, vol. 43, no. 10, pp. 2572–2582, 2011.
- [27] H. Bülow-Hübe, "Daylight in glazed office buildings."
- [28] K. Kawamoto, Y. Shimoda, and M. Mizuno, "Energy saving potential of office equipment power management," *Energy and Buildings*, vol. 36, no. 9, pp. 915–923, 2004.
- [29] C. A. Webber, J. A. Roberson, M. C. McWhinney, R. E. Brown, M. J. Pinckard, and J. F. Busch, "After-hours power status of office equipment in the usa," *Energy*, vol. 31, no. 14, pp. 2823–2838, 2006.
- [30] A. Boyano, P. Hernandez, and O. Wolf, "Energy demands and potential savings in european office buildings: Case studies based on energypplus simulations," *Energy and Buildings*, vol. 65, pp. 19–28, 2013.
- [31] Y. Agarwal, B. Balaji, R. Gupta, J. Lyles, M. Wei, and T. Weng, "Occupancy-driven energy management for smart building automation," in *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*. ACM, 2010, pp. 1–6.
- [32] "The idorm project," <http://cswww.essex.ac.uk/iieg/idorm.htm>, 2013.
- [33] P. Davidsson and M. Boman, "Distributed monitoring and control of office buildings by embedded agents," *Infor. Sciences*, vol. 171, no. 4, pp. 293–307, 2005.