Towards Large-scale Monitoring of Commercial Buildings using Wireless Sensor Networks

Bharathan Balaji
University of California, San Diego
bbalaji@cs.ucsd.edu

Abstract

Buildings consume 73% of the total electricity consumption in the US. To improve the energy efficiency of building loads - HVAC, lighting and plug loads - several methodologies have been developed by the research community in the recent years. Researchers use fine-grained monitoring of buildings to determine the opportunities for improvement in efficiency, and provide better actuation mechanisms based on the information gathered. For such fine-grained monitoring of buildings, wireless sensor networks have been popular choice due to their reduced cost of deployment and ability to be retrofitted within existing buildings. Several different protocols and types of sensors have evolved over the years for managing wireless sensors. However, integrating them to provide a common platform for building automation systems remains a challenge. Some of the initial attempts to integrate different sensor network systems and provide standardized interface for application development are presented here. Deployment of wireless sensor network within buildings provide additional challenges in adoption of wireless sensors based building automation solutions. Experiences and guidelines from previous deployments provide valuable insight to make a large-scale deployment manageable. Thus, published guidelines from some of the largest sensor network deployments have been presented as well.

1 Introduction

Buildings consume more than two-thirds (73%) of the total electricity, 12% of the potable water, are responsible for 39% of the carbon emissions and account for about 46% of the total energy use in the US [57]. Even if a fraction of the energy consumed by the buildings can be reduced, it will lead to huge overall impact to the environment. Further, most people spend 90% of their time within buildings [61]. Modern technologies have the potential to provide a comfortable environment to the building occupants and provide energy efficient solutions for building operations. The opportunity to make an impact in building automation has been correctly identified by the research community and some innovative solutions have been invented in the recent years.

Almost every aspect of buildings has been studied, and new inventions continue to be made every day. To analyze the energy consumption within buildings, Agarwal et al. [9] observed the breakdown of electricity consumption by different subsystems within buildings at UC San Diego over a year. They classified the subsystems within buildings based on energy consumption into four categories- HVAC (Heating, Ventilation and Air-Conditioning), plug-loads, lighting and special equipments like servers, lab equipment, etc. In a separate survey, encompassing three countries - USA, UK and Spain - Perez-Lombard et al. [47] classify commercial buildings loads into three categories - HVAC(48%-55%), lighting(17%-33%) and miscellaneous(15%-32%). To reduce the overall energy consumption of buildings, improvements have to be made in each of these categories.

To monitor plug-loads, plug meters have been developed which measure appliance energy consumption, and allow remote management of devices using smartphones and web browsers [29, 39, 59]. Motion sensors and overhead cameras have been used to estimate occupancy within buildings and opportunistically turn-off HVAC zones during periods of inoccupancy [7, 20]. CO_2 sensors have been used to dynamically control the amount of air flow in HVAC systems based on the number of people in a room [50]. The Sensor Andrew [51] project recognized activities within homes using a combination of motion sensors and energy meters to identify wastage of electricity. In order to facilitate rapid deployment, wireless sensor networks have been a popular choice for research within buildings. Further, installation of wireless sensors on a large scale in an existing building is much cheaper compared to wired sensors. However, many challenges remain to be solved to make wireless sensor networks a pragmatic solution.

Deployment of different kind of sensing technology within buildings often employ an assortment of protocols which are incompatible with each other. Wired sensors like temperature, humidity and motion sensors use BACNet or LonWorks protocol in modern buildings. Wireless sensors either use proprietary protocols or one of the various standardized protocols based on 802.15.4 radio - Z-Wave [4], ZigBee [5], 6loWPAN [52]. Further, each sensor generates data that is unique to its sensing modality and implementation. Energy meters send out periodic power values, thermometers send periodic temperature data, occupancy sensors send stochastic messages of a region's occupancy. To manage these diverse types of sensors, analyze the measurement data and develop useful applications on top of the sensors deployed, a centralized management system needs to be developed. HomeOS, Sensor Andrew, sMAP and Building Depot are initial approaches to develop such a centralized system [8, 16, 18, 51].

Another challenge for commercialization of building monitoring system is the deployment of sensors in modern buildings. Wired sensors need to be installed in the building when it is under construction. Retrofitting wired sensors is prohibitively expensive and aesthetically unpleasant. Wireless sensors, on the other hand, face a different set of challenges. The RF environment within a building calls for careful placement of sensors, with certain amount of overprovisioning to compensate for weak links. Furthermore, configuration of sensors is a manual process. Sensors need to be configured with their type, the location within the building, the users which have access to it, etc. To scale to thousands of wireless sensors, techniques need to be developed to configure each sensor with a minimum amount of effort. Experience from preliminary deployments also reveal that user interference can lead to unexpected down times for sensors, and location within a room and the aesthetics of the sensors contribute significantly to the acceptance of the system by the users [28].

This paper surveys prior work which has addressed the large scale deployment of sensors within buildings addressing each of the aspects described above. The rest of the paper is organized as follows - Section 2 provides an overall view of the sensing technologies available today for buildings, Section 3 discusses the issue of integrating the various sensing solutions in to a single modular system, Section 4 provides an overview of standardized programming interfaces being developed for application development in a building automation system, Section 5 studies the challenges involved in large scale deployment of sensors and summarizes lessons learnt from previous deployments, Section 6 provides a brief discussion and future directions of research. Section 7 concludes the paper.

2 Sensors

Several innovative ideas have been developed over the years to monitor different aspects of modern buildings, and provide insight in to the characteristics of daily operation. This section presents some of the key papers addressing various avenues of research in this area. The papers discussed are not intended to be an exhaustive coverage of all the research related to building sensor networks, but rather chosen to bring out the different types of sensing modalities required in each type of monitoring application.

2.1 Energy Meter

Energy metering of buildings is a well studied subject, with the first power meters developed as early as 1889. They have been used in daily life for measurement of building wide consumption, or on a per aparment basis for billing purposes. However, in order to identify opportunities to conserve electricity consumption, a more fine-grained instrumentation is required.

2.1.1 Non-Intrusive Load Monitoring

In 1992, Hart invented a technique called non-intrusive load monitoring (NILM) [27]. The power traces obtained from the power meter for the entire house is analyzed to glean information about the power consumption patterns of individual appliances. Each type of appliance has a unique power signature, with each appliance identified by a combination of its average power consumption, different states of operation and power factor. Further, appliances are switched on and off at different times, showing a distinct state change in the total power trace as shown in Figure 1.

Models are developed apriori for every appliance class based on their states of operation like, off-low-medium-high for a table fan; power signatures during turning on, steady state, transient and powering down operations; frequency scatter from the fundamental frequency of operation (120V, 60GHz); and finally, harmonic frequency signatures. Several instrusive techniques are also explored for initial collection of data to train and perfect the model. The algorithm for disaggregation from the total power trace employs edge detection, cluster analysis and maps the results with the models developed. The loads are then tracked to give a power breakdown by every appliance within the home.

There are several limitations to the NILM approach. First, a model needs to be developed for every class

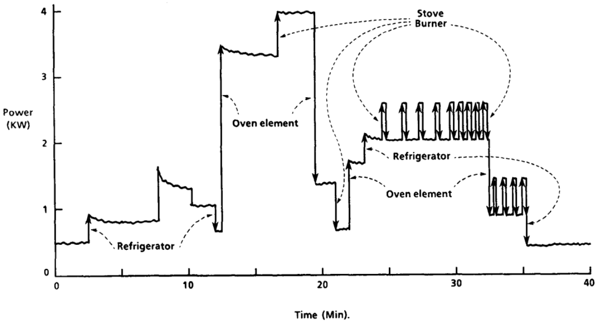

Figure 1: Total power consumption in a home for a period of 40 mins. The trace obtained during culinary activities within the residence, leading to state changes by various appliances [27]

of appliance in a building. This list is exhaustive, and hence, the task is labor intensive, if not prohibitive. Further, only a few state changes of different class of devices can be tolerated by the system at a time. Several improvements [36, 23, 49, 37] and improvisations [45, 33] have been made in the NILM algorithms since Hart published his work.

In a home environment, with a few people utilizing different kinds of appliances at different times, NILM provides an unobtrusive method for fine-grained measurement of power. However, in a commercial building environment, many instances of the same class of appliances are used, and several state changes are likely to happen at the same time. Another limitation has developed over the years with the prevalence of highly efficient switching power supplies for various appliances, which the NILM algorithms find hard to distinguish.

2.1.2 Plug load meters

Another approach towards fine grained measurement of electricity is plug load meters for individual loads in the building. MIT Plug [39] was one of the first approaches to address this problem. They developed a power strip which could measure power, light intensity and audio using appropriate sensors. The design was a simple one, utilizing a proprietary wireless protocol for communication. They deployed 27 of the prototypes in their lab and showed it was possible to monitor individual appliances and draw inferences from them. ACme [29] was developed by UC Berkeley two years later. Careful design ensured a much more accurate measurement of the plug load, with the addition of the actuation facility for the first time. The meter uses Epic microcontroller running TinyOS, with an ASIC ADE7753 for energy measurement. The wireless communication was through 6loWPAN, the embedded version IPv6 protocol, running on top of 802.15.4. Various parameters - voltage, current, active power, power factor could be measured. Data was transmitted at once per minute. ACme proved to be a standard design for plug load energy meters, with several research groups using similar designs for their projects [31, 17]. Figure 2 shows the internal design of ACme.

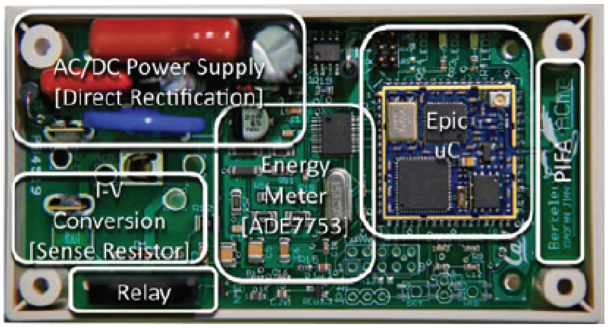

Figure 2: ACme Energy Meter[29].

Synergy Energy Meter[59], developed at UC San Diego, was an additional improvement over the ACme design. The design used a simpler microcontroller, and a system-on-chip radio, reducing the cost of the meter significantly. Further, reliable actuation capability was provided using a mechanical relay. However, the unique aspect of the design was that the meter could support software controlled configurations for metering and actuation. The data rate could be throttled to transmit only when interesting events happened, devices could be set to different priorities and higher level applications could exploit them for different mechanisms. Several demand response policies were showcased using the mechanisms provided.

To comprehensively measure all the plug loads within the building, plug meters would have to be installed for each and every appliance in the building. Deploying such a vast array of wireless meters takes a lot initial investment and effort. Initial such deployments [59, 17] indicate that the majority of the power is from personal computers. Unfortunately, computers cannot be switched off from the plug directly. Many appliances like table lamp and cell phone chargers do not draw as much power, and meters need not be added for them. An optimal installation of meters targetted towards the power heavy appliances within buildings, to extract the maximum energy benefits with minimum effort, remains to be done.

2.2 Occupancy

Occupancy has long been an object of interest in research - both in homes and commercial spaces. Occupancy en-

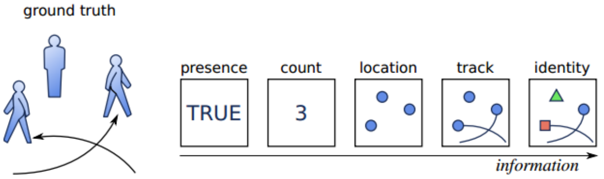

Figure 3: The five spatio-temporal properties of human sensing[55].

ables huge number of applications, impacting various disciplines: health care, building safety, home automation, traffic monitoring, etc. Tiexeira et al. [55] provide a good survey of different types of human detection system as of 2010. They categorize the human sensing system into four major types: presence, count, location, track, identity. They are illustrated in Figure 3. Each type of sensing offers a different level of granularity, and enables different types of applications within buildlings.

Various avenues of detection has been explored using sensing human traits, including gender identification using human odour [46], detection using CO_2 [56], tracking [42] and identification [43] using pressure sensors, electric field sensors [58], and tracking using acoustic sensors [48]. However, these sensors either need a dense deployment of sensors, or do not provide high accuracy in a real environment to scale to large deployments in buildings.

Scanning range-finders using radio waves (radar), sound/ultrasound (sonar), visible light (lidar) and laser (ladar) try to create 2D or 3D snapshots of the environment by transmitting a signal and measuring the response echo. They tend to be expensive and power consuming as they require multiple transmitters in a phased array. They also tend to be noisy in indoor environments due to multipath and scattering effects. The latter problem was overcome by Zetik et al. [62] using an approach common in computer vision, called background subtraction.

Passive Infra-Red (PIR) sensors are used for sensing motion to control lighting in many modern buildings today. However, PIR sensors are not accurate enough for other applications which require precise occupancy information, like HVAC or appliance control. The main disadvantage of PIR sensors is they cannot detect people who are stationary, thus leading to a large number of false negatives and their output is highly bursty. A combination of PIR sensors and door sensors have shown to be much more accurate for single room offices in prior work [7, 40]. Agarwal et al. [6] used these motion sensors in a real environment for controlling HVAC system in a building at UC San Diego.

Simple doppler-shift sensors are also used as motion sensors in today's environment, similar to PIR sensors. However, they have been shown to detect *stationary* people from the motion of their breathing lungs [22]. Doppler radars have also been used to differentiate between 0 person, 1 person and more than one person by detecting heartbeats [63].

Cameras provide an affordable sensing mechanism providing a gamut of information in a scene, including size, shape, color and so on. A popular method for people detection is background subtraction, assuming that the background is either static or slowly changing [38, 10]. Thermal imagers can better differentiate people from background objects through their temperature, but the cost of the equipment has discouraged widespread use [1]. Several *smart cameras* have been developed that extract motion information at the hardware level [26, 30]. Erickson et al. [20] used such a camera system for estimating occupancy within a university building, and estimated annual savings of 42% on the HVAC energy consumption. Commercial solutions have also used cameras as virtual turnstiles to count people coming in and out of an entrance [3].

Despite the plethora of research in human detection and tracking, few solutions have been adopted in practice. Motion sensors and RFID badge based security system are common in many of the enterprise buildings today. However, tracking individual occupants, or detecting their presence in each room of the building has not been adopted due to low accuracy in realistic environment, or high cost of installation and maintenance. Occupancy information plays a key role in actuation of building loads, and thus, a solution which would provide high accuracy without incurring high cost is a stepping stone to reduce building energy consumption.

2.3 Water

Water is another resource which is used extensively in buildings. A sensing system which can detect leakages and provide feedback of usage to the residents will be invaluable. Several approaches have been tried to implement a centralized flow meter, similar to the NILM approach for power measurements [32, 25, 21, 24].

Of all the approaches, Hydrosense [25] is the most recent approach for monitoring water usage on a per valve basis in homes. It uses a single point of sensing the water flow, similar to the non-intrusive technique used for disaggregating electricity consumption.

The plumbing system in residences maintains a constant pressure throughout the piping when no water is flowing. The instant a valve is opened or closed , a pressure change occurs and a pressure wave is generated in the plumbing system. Transient pressure wave phenomenon results from the rapid change of water velocity in a pipeline, called as a *water hammer*. The magnitude of the pressure change is positive or negative depending whether the valve is being opened or closed. Further, the water hammer signature for a particular plumbing system depends on the valve type and the location in the pipe network.

A high quality pressure sensor (0.002 psi) sampled at a rate of 1kHz at any exposed pipe within the home can be used to distinguish the occurance of valve open/close events as well as the flow rate, when Poiseuille's Law is applied. The fixtures are identified by a combination of the signal detection techniques - thresholding and monitoring rate of change of pressure to detect valve events; matched derivative filter to classify among fixtures. Hydrosense also requires an initial calibration of every valve for flow estimation.

Froehlich et. al. developed a prototype and deployed it in ten residences and evaluated the accuracy of the system to be 97.1%. The accuracy of the system decreased when a single pressure regulator was shared between multiple homes. Another feature of the system is that it can distinguist between events which overlap with each other, although it cannot distinguish between events occurring exactly at the same time. The system will scale well for isolated residences, but the accuracy is likely to drop in an apartment based residence.

For a commercial building or for an apartment complex, the system is unviable in its current state as the flow rate estimation algorithm does not take the changes in pressure or gravity into account, which may be significant in long vertical pipes. Callibration and classification of more fixtures would also be more difficult. Installation of more points of sensing, with each sensor identifying a portion of the building(say, an apartment) would help to increase the accuracy of the system in large buildings.

2.4 Others

Several other types of sensing technologies have been developed and many of them are already deployed in modern buildings. Thermostats in buildings today measure the temperature and humidity of the room, provide limited control to user for setting the temperature of the room and provide a means to indicate user occupancy during off hours. Photo-sensors are used to control the brightness levels of lighting in rooms which can receive sun light during the day. HVAC components use a combination of sensing technologies - air flow meter, temperature sensors, heating output of individual zones, etc. Commercial solutions exist which use CO_2 sensors control air flow from HVAC system based on the number of people in the room [50].

Each type of sensor described in this section has a different set of requirements - varying data rates, periodic or aperiodic events, specific location within the building for optimal operation, etc. Further, each type of sensor may choose to use a different wireless/wired protocol depending on the device, vendor and building manager. A framework is required to capture the data generated by all these types of sensors for ease of installation, use and maintenance by the building personnel. Such a framework needs to solve challenges in not only deployment, maintenance and collection of data generated by the sensors, but also needs to address privacy and security concerns, and ease of use for the building manager. In the following sections, we will see papers which have addressed some of the challenges associated with a sensor network based building management system.

3 Building Management Framework

The goal of the building management framework is to provide an organized structure to not only the sensor data from various sources, but also to manage the data collected and provide the information required to the different control subsystems like HVAC and lighting. Such a framework would have to consider each and every aspect of the building it is deployed in. We look at three different approaches to solve this problem - the HomeOS [18], the Sensor Andrew project [51] and the Smart Grid Information Model [64].

3.1 HomeOS

Dixon et al. [18] have developed an operating system for several types of applications in a home environment. The operating system is designed by abstracting home automation technology as a Personal Computer(PC). The HomeOS addresses three objectives. First, they provide a central management platform for all the devices for a non-expert user. Second, they provide a platform for application development which can be easily configured in different kinds of home environment, devices and user control. Third, they make it easy to add a new device to the system and prevent vendor locking.

The HomeOS consists of four layers as shown in Figure 4. The *Device Connectivity Layer* (DCL) is responsible for providing a common interface across different types of network protocols used by appliances. There is a module for every protocol, like DLNA, Z-Wave, ZigBee. The *Device Functionality Layer* (DFL) is responsible for providing handles for different kinds of devices. DFL uses the DCL APIs and abstracts the functionality provided by every appliance, like TV, energy meter, lights, etc.

The *Management Layer* provides a central place to add and remove applications, devices and users as well as set policies for access control. Every user and application need to have the necessary permissions to access

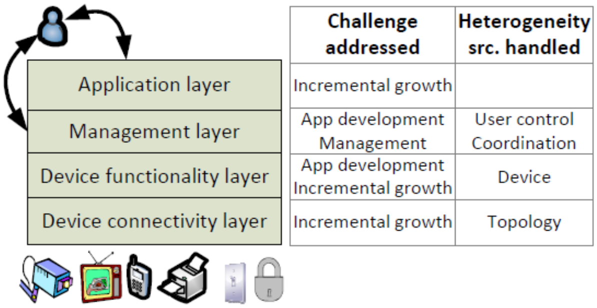

	Challenge addressed	Heterogeneity src. handled
Application layer	Incremental growth	
Management layer	App development Management	User control Coordination
Device functionality layer	App development Incremental growth	Device
Device connectivity layer	Incremental growth	Topology

Figure 4: Layers in HomeOS and their considerations [18].

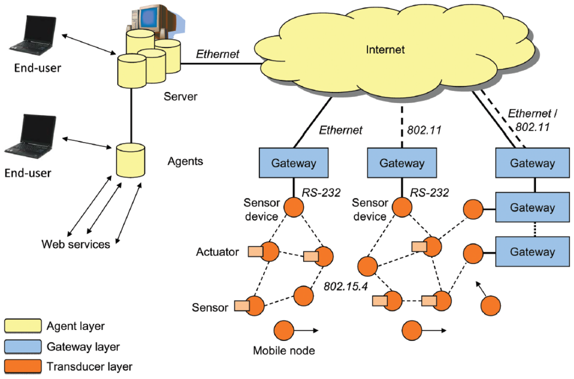

Figure 5: Three-tiered architecture of Sensor Andrew. Dashed and solid lines represent wireless and wired connections, respectively. Actuators and sensors are both transducers and can be attached to sensor devices or mobile nodes [51].

other devices. The access control policies are setup as Datalog[15] rules. The *Application Layer* provides the environment for developing applications across the devices in the home. Each application has to provide a manifest describing what services they need, similar to the Google Android [19].

HomeOS is a great initiative to integrate different types of sensors, and provide an interface to the occupant that is easy to use. However, the framework is meant for homes, and scaling issues that come up with an enterprise building are not addressed. For example, the access control rules using Datalog is complex within homes, which each occupant having different privileges for every appliance depending on time of day. Further, guests are given a different set of privileges. For an enterprise building, the number of rules will grow exponentially higher. A hierarchy of rules need to be introduced, classifying depending on the type of the occupant - staff, janitor, building manager, etc. Similar issues exist in application development and personalized control as well.

3.2 Sensor Andrew

The Sensor Andrew[51] system developed in Carnegie Mellon University addresses the problems similar to the HomeOS in a commercial building environment, on a university-wide scale. The architecture is built around the Extensible Messaging and Presence Protocol (XMPP). It is a standard scalable messaging and presense protocol with user/group authorization, authentication and access control. It provides a publish-subscribe messaging service in a client-server model. This is the basic model used by Sensor Andrew for communication among the sensor devices, users and applications.

Several components are added to the system around the basic XMPP transportation model as shown in Figure 5. The *Transducer Layer* provides adapters to Internet-connected devices from all the end-point sensors and actuators. It is responsible for supporting all the different types of devices and communication protocols. The *Gateway Layer* consist of devices which have access to the Internet and are configured as XMPP clients. The gateways collect all the information from the transducer layer using low-level protocols and passes it on the *Server Layer* using XMPP through event nodes. Devices in the server layer need to subscribe to event nodes to receive published data. The server layer supports all the client applications on the XMPP servers. A web application, called Data Handler, is used to maintain the schema, business rule and read/write functions. It provides the interface for browsing, editing, creating transducer and device metadata records in the registry. Security is provided through access control lists.

Rowe et al. demonstrate the capability of the Sensor Andrew system in a home environment. Energy meters and motion sensors are deployed throughout a home, and are used to infer the appliances which show high correlation between motion and energy use. Energy wastage is identified when the appliances which show high correlation continue to use energy without accompanying motion in its surroundings. The framework is also used to detect anomalies in energy consumption when the energy consumption deviate from the average consumption values.

The Sensor Andrew system provides a good communication framework among sensors deployed in a large scale deployment. However, an evaluation of the system in a real deployment has not been shown in the literature. The applications show-cased were developed in a home setting. The system itself does not address the ease of application development and use by building managers. There is no structure given to the data obtained from different sources, for querying status of sensors, ease of de-

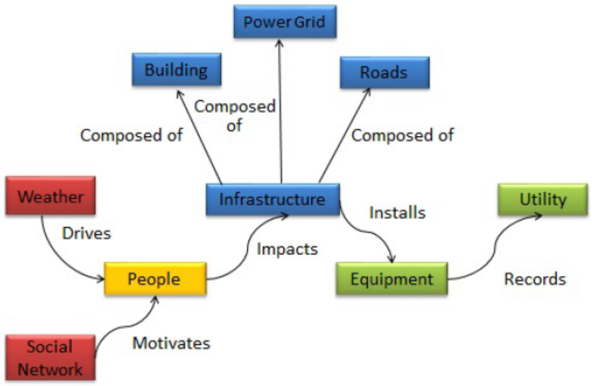

Figure 6: Interplay between information concept spaces that are interesting to Smart Grid Applications represented using Web Ontology Language [64].

velopment and maintenance.

3.3 Smart Grid Information Model

Simmhan et al.[53] and Zhou et al.[64] have built a semantic model for buildings in University of South California. The purpose of the system is accurate demand forecasting and effective load curtailment during a Demand Response (DR) event in a Smart Grid environment. Unlike Sensor Andrew and HomeOS, the semantic model does not seek to provide ease of management for fine grained control of sensors/actuators within the building. It uses the information collected from the sensors to suggest strategies to reduce energy consumption during a DR event using analytics.

The structure of information collection is similar to the Sensor Andrew system. A *Transport Agent* acquires raw data from remote sources - sensors and actuators - using different access protocols. A *Parser Agent* extracts specific data attributes from the raw data. A *Semantic Agent* annotates and maps the data, attribute tuples into meaningful concepts that are described by a semantic ontology. The data is collected in a Jena database, and machine-learning techniques are used to build forecasting models. Events processing is used to find non-intuitive patterns in the data which might help save energy.

The unique aspect of the framework is that it uses Semantic Web[2] to model the building information. The semantic model uses Web Ontology Language(OWL) to represent the knowledge about each aspect of the system. The *Electrical Equipments Ontology* captures the information pertaining to power distribution system, the equipments installed as well as their measurement units and categories of equipments - like, light, vending machine, sensor, etc. The *Organization Ontology* is used to represent the organization within a company - the director, facilities manager, human resource manager and so on. *Infrastructure Ontology* captures the information about the building structure, surrounding transportation network, type of the building, etc. *Weather Ontology* represents the long term and short term weather patterns as well as natural phenomenon occuring at a location. *Spatial Ontology* captures properties about the location - latitude, longitude, altitude, zip code. *Temporal Ontology* tries to capture information that happens on a regular basis, or the events which have been scheduled on a calendar. The integrated ontology forms the relationships between all the component ontologies. The ontology schema as well as the instance data is stored in MySQL database and querying is performed using SPARQL.

The Smart Grid Information Model provides a framework for organizing all the data about buildings. Although it is not meant for real-time control of buildings, integrating the information model with a system like Sensor Andrew or HomeOS would help in organizing the data from diverse sources, analyzing the data for usage patterns and ease development of applications.

4 Application Programming Interface

We have seen some solutions which integrate different parts of the building automation system. An emerging pattern is a layered structure, every layer addressing challenges at each level of abstraction. An important aspect of the integration all these layers is played by standardized protocols and interfaces, providing a means for faster development across different types of vendors and systems. One of the popular standard protocols, BACnet, and recent work in development of APIs are presented here.

4.1 BACnet

American Society of Heating, Refrigerating and Airconditioning Engineers (ASHRAE) is responsible for many of the building standards and guidelines in use today. The society recognised the need for a distributed control system that provide HVAC control, lighting control, security, fire detection/suppression and other building systems. BACnet was introduced in 1995 [13] as a standard communication protocol for building automation and control network.

Figure 7 shows the basic architecture of BACnet as per the 1995 standard. BACnet application layer provides a model of information contained in a building automation device and provides a group of services to exchange that information. BACnet objects are used to abstract the information of the devices using standard object oriented design practices. Each BACnet device is restricted

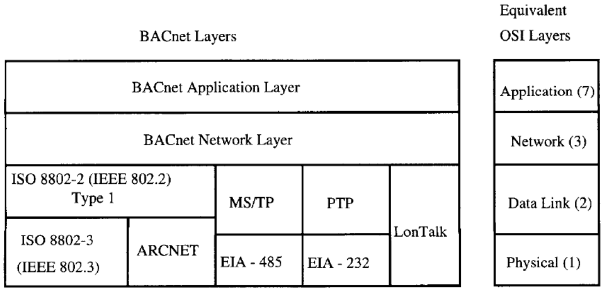

Figure 7: BACnet collapsed architecture [13].

to have only one instance of "Device" object, and the "Object_Identifier" field uniquely identifies each device in the BACnet network.

The *Alarm and Event services* in the BACnet application layer provide a way to subscribe to change of value notifications, request a status summary for alarms or events, notify devices that alarms or events have occurred, and acknowledge that an operator has seen an alarm notification. *File Access services* provide the means to read and write files atomically, including the ability to upload and download control programs and databases. *Object Access services* provide a means to read, write, create and delete the properties of the objects. *Remote Device Management services* provide tools for troubleshooting and maintaining devices. *Virtual Terminal* provides a console for user interaction. *RequestKey and Authenticate services* provide security features.

BACnet network layer does not support all the functionality of the equivalent OSI Reference Model. Atmost one path can exist between two devices. Message segmentation and reassembly is not supported at the network layer, and the maximum length of the packet should not exceed the capability of any of the data link technology encountered in the path from source to destination. If destination for the message is on the same network, no additional network layer information is needed. If the destination is on a remote network, the client device must include the destination network number and MAC address of the destination device. The router on the local network will insert addressing information about the local network, so a device does not need to know its own network number. Support for IP was also provided using tunnelling through Packet Assembler Disassemblers (PADs), which works similar to a gateway.

At the time of release, Ethernet was the most popular option for the data link layer. Other alternatives were also provided to support other technologies popular at the time - ARCNET, a token passing protocol; EIA 485, a building control system physical layer technology along with Master-Slave/Token Passing for medium access control; LonTalk, a patented, propreitary protocol developed by Echelon Corporation; and, Point-to-Point protocol.

Several additions were made to the second version of the standard, released in 2001 [14]. BACnet/IP support was added to the networking layer, allowing BACnet devices to communicate directly with each other over the Internet. Additional device profiles were added to support interoperability between different manufacturers. Some of the restrictions from the 1995 standard were removed, allowing for easier development and remote management of devices. Trending supported was added, allowing data logging of measured device quantities. Several extensions were also added, so that BACnet can be used for automation of other building services - life safety systems, lighting control applications, access control systems, security within the BACnet framework and methodologies of test for conformance with BACnet.

Since its release, BACnet has been widely adopted all over the world for building automation. Other competing protocols exist - LonWorks, ModBus, OLE for Process Control (OPC). However, BACnet has gained the majority of the market share as it is an open standard, covering a wide spectrum of issues in building automation and is easily extendable for each type of installation. WSN devices can be added as a natural extension to the BACnet by incorporating the relevant protocols in to the network and the data link layers. The ZigBee Alliance is working with the BACnet committee in this direction [41]. Park et al. [44] have provided an implementation based on the ZigBee 2004 specification.

4.2 sMAP

sMAP - a Simple Measurement and Actuation Profile for Physical information, was developed by Dawson-Haggerty et al. [16] as a common interface for all types sensors and actuators. Although not specifically meant for building sensors, it has been implemented as part of several building projects successfully [17, 34]. The API design is divided into three areas: Metrology, Syndication and Scalability.

Metrology: Each sensor is assigned a measurement point specifying its location and a channel to designate the type of physical quantity being measured. Only scalar units are supported, and each measurement is associated with a timestamp, sequence numbers, a scaling coefficient and a measurement unit. Common modalities of actuation are also supported.

Syndication: This part of the design concerns with the propogation of the data into a larger system. The sensor data and metadata are exposed using HTTP following the Representational State Transfer (REST) paradigm. Data is represented using JSON objects and structured using

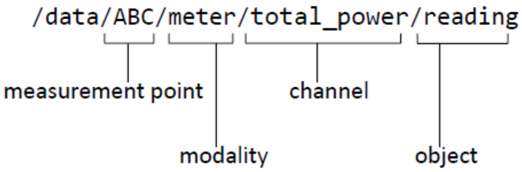

Figure 8: Canonical sMAP URL for the "total power" meter of a three-phase electric meter [16].

information schema. An example of a URI is shown in Figure 8. As HTTP supports a client-server model, it encourages polling of data rather than a push based data collection system. However, most sensing data are event driven or generate periodic reports. To support such sensors, a `/reporting` resource is created, which allow to specify client URLs for reporting the data.

Scalability: sMAP provides a framework for the sensing devices to host the API interface. The networking layers are supported by 802.15.4 (MAC), 6lowPAN (IPv6)and HYDRO (Routing) with TCP and UDP in a TinyOS system. A compact form of HTTP, called Embedded Binary HTTP (EBHTTP) is used, which is obtained by binary-formatted, stateless encoding of the standard protocol. JSON objects are also compressed to binary format, called packed JSON. The combination of all the adaptations allow internet supported sensing devices to directly host sMAP even with limited computing resources.

sMAP provides a solid abstraction of all types of sensors that can be found in a building, and exposes them to other layers using REST APIs. It can connect the device/transducer layer in a building management framework to the rest of the system. However, as a simple API, it does not address any of the concerns of the layers above or below the interface.

4.3 Building Depot

An important part of the system is the storage, access and sharing of data from the building sensors. Agarwal et al. have designed an API, called Building Depot [8], based on a distributed architecture for handling data storage. The API uses RESTful HTTP service, similar to the ideology behind sMAP, and uses JSON objects for data representation. Figure 9 provides an architectural overview of Building Depot.

The central actor in the system is a resource, which is the data that the client desires from the server. A uniform set of verbs are used to interact with the representations of that resource. For example, a resource might be occupancy in a given room, and the representation of that would be an JSON message that describes the occupancy. *Data Connectors* are developed for each of the protocols, so that all the sensors using the same protocol can reuse the connectors. *Data Service* is the module which handles all the sensor data and makes it available to different applications. The module also stores contextual data, like location, type of data, access control lists for each sensor.

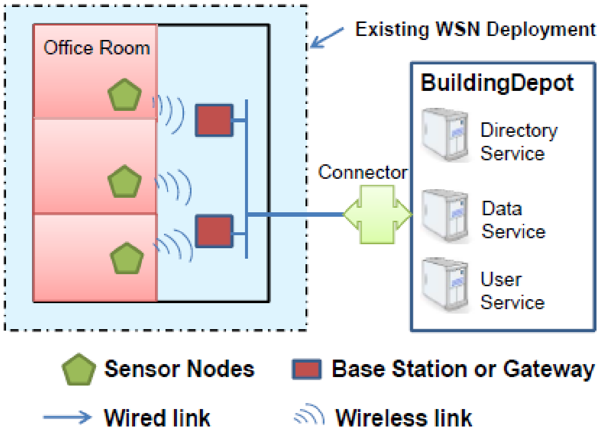

Figure 9: High-level view of BuildingDepot. BuildingDepot consists of the three services. Each institution has one top-level Directory Service, one User Service, and any number of Data Services [8].

The *Directory Service* contains links to the various Data Services and the Directory Services beneath it. These Directory Services are used to form a hierarchical tree of all of the Data Services and Directory Services that make up an institutions Building Depot system. The *User Service* provides user management to ease data sharing and administration. Users are authenticated using a registration system similar to most personalized systems in the internet today. The *Community Service* is a simple directory of all public Building Depot deployments and is meant to facilitate directory lookups and sharing across the world.

The APIs are designed to encourage development of applications across the data collected by the sensors. Some of the applications the authors envision being implemented are data dashboard, building visualizer and occupancy-based actuation.

Building Depot serves a building management system akin to how a hard drive serves an operating system. It provides a management framework to store the vast amount of data collected from different sources in an organized manner, and provides a uniform interface to access this data. The drawback of the Building Depot system is that it has not been evaluated in a realistic environment. It remains to be seen if it can fulfill the promise of a standardized interface between storage and rest of the building management system.

5 Deployment

Deployment of wireless sensors within buildings have completely different logistics than the conventional wired sensors. In this section, we will explore the deployment experiences that have been published by several research groups over the years.

5.1 Deployment in Swiss Alps

Wireless Sensor Networks (WSN) is a relatively new field, and deployments in real world were not attempted until 2002[54]. Many deployments have been attempted since then, including one on top of an active volcano[60] and a failed attempt in a potato field[35]. Barrenetxea et. al. [11] have provided a hitchiker's guide to deploy wireless sensor networks, and they cover a wide range of aspects for successful, scalable, maintainable deployments. The system under test was a habitat monitoring system for the Swiss Alps to predict avalanches and mud streams. The salient features of the guide are presented here briefly.

Local conditions like temperature and humidity can have unexpected effects on the electronics. Although the weather conditions are not expected to be extreme within buildings, the electronics can be exposed to extreme conditions depending on its location - near hot water pipes, or damp area in the kitchen. Further, wireless connectivity are known to be hard to estimate and time variant within buildings, with the range varying with time and location. The hardware should be developed and tested for all such conditions.

The electronics should have a proper encasing, following all the safety certification requirements. A poor case can lead to damage, or worse - incorrect readings, over time with exposure to dust, sunlight and water.

Once deployed, the sensors can be difficult to access. Remote monitoring, resetting and reprogramming of the sensor nodes are must have features in the system. Remote monitoring should be done with periodic logs about various aspects of the system - battery life, network link status, sensor readings.

Testbeds offer important insights in to the flaws of the system. However, testbeds themselves should be carefully designed to reduce wasted effort and offer flexibility. The motes should be easily reprogrammable, easily accessable and preferrably powered by AC mains to avoid problems because of battery life. Further, no changes should be made to the tested sensors before they are deployed in the real world. Last minute changes, especially in software, can cause unexpected bugs which can be difficult to debug.

Clear labels on the motes, and informing the occupants in the area by others means about the testbed reduces the chances of disconnection of the device by people.

The sensor motes are generally caliberated for accurate measurement readings. Data obtained from the sensors should be verified with a reference instrument for a period of several days in both the laboratory and the deployed settings. In case of re-deployment, calibration should be again verified.

Simulations go a long way in filling the gap between theory and real world deployements. Build and simulate models to test the protocols and common use cases.

5.2 Home Deployments in Virginia

Hnat et al. [28] provide a set of guidelines based on the experiences of deployments near University of Virginia covering 20 homes over the course of several years. Various types of sensors were deployed to study different aspects of home monitoring (over 1200, in total). The sensors included motion sensors from X10, Aeon Labs Z-wave contact switches, GE Z-wave home automation switches, GoMotion ultrasonic range finders, TED 5000 for whole house metering, Powerhouse Dynamic eMonitor, Aeon Labs Z-wave meter for plug load monitoring, Shenitech's ultrasonic flow meter and platforms for monitoring light, temperature and humidity levels.

A gateway was setup to convey the data from each of the sensing systems to the central database via the internet. Each subsystem came from different vendors, using their own gateways for collecting information from the sensor nodes. Some of the sensor nodes relied on the AC power, and some were completely battery powered. This led to varying levels of data loss during power outages, end of battery life and broadband disconnections. The root causes of data loss over a seven month period of four homes has been shown in Figure 10. Some of the salient points in their hitchhiker's guide are discussed below.

Despite the apparent abundance of power, the authors found it difficult to install and maintain sensors which use power out of wall sockets, especially when the number of sensors began to scale. The wiring required with limited number of receptacles and sensors which needed to be accurately positioned increased both installation and maintenance costs. The residents routinely detached sensors to access the power sockets for temporary use. Further, the wires got damaged accidentally by children, pets and robotic vacuums. The authors found that battery failures were far easier to predict and debug in most cases.

Some of the sensors used power from the main power supply and were found to be more reliable than using receptacles. However, for any maintenance, the electricity for the entire home needed to cut off. The major cause of data failure in these sensors were found to be power outages. Experiments with indoor solar power found that

Root Cause	House G	House H	House I	House J
Sensing Sub-system	4107	642	4757	274
Gateway Down	5596	0	3	136
Plug Disconnected	509	30	474	10
Battery Dead	452	17	168	0
Wireless Link Loss	410	0	122	1
Internet Down	251	97	178	9
Power Outage	21	0	87	2

Figure 10: The total sensor down time for four home deployments over a seven month period, broken down by root cause and measured in sensor-days: #days * #sensors [28].

the power was inadequate and did not correlate with the active periods in homes.

Wireless connectivity can be poor in home environments despite the small area to be covered. Sensors are often placed in metal boxes, and placed in exceptional locations. Attenuation occurs due to plaster, masonry, concrete and heavy metal appliances. Further, as the number of sensor subsystems increase, the available bandwidth has to be shared using frequency multiplexing.

As the number of sensors scale, the maintenance required increases considerably. Children get attracted to wires and LEDs, guests knock off sensors accidentally or get knocked off by vacuum cleaners. The authors suggest automatic component-level checks and end-to-end data verification. Some of the parameters to be checked are: network connectivity, service availability, last entry time, calibration, time stamps check and CPU load on end hosts.

Deployments in home usually has to be done within a minimum amount of time, to reduce the inconvenience caused to home owners. The authors suggest to break the deployment to make this possible - for site evaluation and for deployment itself. To reduce the on-site time, all the background work and planning of deployment needs to be done beforehand.

User participation is as limited as energy in a battery. The authors found that users can report high quality, fine grained information for short period of time (few hours), and can provide coarse granularity information with increased period of participation (few weeks). This is not due to lack of motivation, but because people can tolerate only so much intereference due to participation. The authors suggest using several sensing modalities to verify the ground truth instead.

Aesthetics become an important consideration as the number of sensors within home scale to hundreds. The authors suggest minimizing exposed wiring, mounting the sensors in such a manner that they are noticeable in everyday routine. Further, any LEDs or sound making sensing nodes should be avoided, as users tend to get annoyed with them over a period of time.

5.3 Building Deployment at Berkeley

Dawson-Haggerty et al. [17] have conducted a long term deployment study in a commercial building at Berkeley. The deployment consisted of 455 meters across 4 floors with a staged stratified sampling method to analyze Miscellaneous Electric Loads (MELs) in the building. The meters were based on the ACme energy meters developed by Jiang et al. [29].

In the initial stages of deployment, a survey of all the MELs in the building was conducted and a taxonomy was developed based on End Use, Category and Product Type. To provide stability to the wireless network, Load Balancing Routers (LBRs) were installed at strategic locations to act as routers for the IPv6 enabled energy meters. The LBRs were also used to facilitate remote debugging. They had overlapping regions to increase reliability in case a LBR breaks down. The authors claim that using IPv6 allowed them to develop compact implementations of services and made it easy to scale to large numbers.

The software on the meters were designed such that simple configuration parameters like sampling rate and calibration parameters can be extracted and changed on the fly. The authors found this utility extremely useful, and had to seldom resort to re-programming the devices using over-the-air image update. Time synchronization was established using the Global Time timestamps. Sequence numbers were used to analyze network performance and local time using 32kHz crystals were used within the meters to detect any network outages or device resets.

The authors found that data collection (mutipoint-to-point)and dissemination (point-to-multipoint) are the dominant traffic patterns, and point-to-point communication is rarely needed. Another interesting observation was that data loss occurred more due to device unpluggings by users than due to network connectivity. If perfect reliability is required, the authors recommend buffering at each layer which can fail.

Although the meters were static within the buildings, the network links were found to be dynamic. Of the total of 1200 links, only 500 were static and rest of the links changed as the stability of the links changed with time. The authors believe these dynamic links are due to the noise floor and interference conditions, since the network remains fairly static over weekends.

6 Discussion and Future Work

The overall objective of the building management system is to provide a comfortable environment for the occupants of the building, spending the least amount of energy. A more ambitious goal is to have an energy proportional building, where the energy consumption is negligible when the building is not in active use. Further, the building management system has the capability to provide a personalized environment to each occupant.

Although prior research efforts have taken great strides much towards accomplishing this goal, we are still a long way from fully achieving it. As is evident from the different research efforts we have seen so far, building management systems are complex systems with various types of subsystems working together. A common management framework for all the sensors within a building needs to address issues that goes beyond a single field of study.

Humans have almost been completely left out of the loop in the systems being developed. The comfort level of the occupant can only be measured by insightful surveys, providing valuable feedback to the management system. It is evident on how minor things can effect the usability of the system by the deployment experiences we discussed in Section 5. All the different aspects of buildings - lighting, HVAC, appliances, water - affect the occupants within the building in some manner. It is critical to understand the feedback provided by the users of the building to develop a building management system that boosts the productivity and comfort level of the occupants.

Security and privacy of the occupants present a different set of challenges. For example, recent work has shown that smart meters installed in a home can reveal information about type of ownership of the home, number of appliances, number of occupants and employment status of the occupants [12]. Addressing such issues is key to wide-spread adoption of solutions which require fine-grained monitoring of a building. The system should be able to provide a trust framework in which the users are comfortable sharing their personal data without sacrificing their privacy.

Several systems being developed encourage ease of development of applications on top of the building management system. However, little effort has gone towards ensuring that the applications do not violate any of the security or safety codes within the buildings. Further, the application itself should not interfere with other applications which might be critical for normal building applications. The building management system has to provide an environment equivalent to a sandbox for safe application deployment.

The above issues are just some of the challenges faced by the developers of the next generation of building management system. There are implications in several areas of computer science - wireless sensor networking, embedded systems, distributed systems, database management, security and human computer interaction. As humans form the center of the system, non-engineering fields such as socialogy, ethnography and psychology also play a crucial in understanding the needs of the building occupants. It would be interesting to see how the building management systems evolve which take in to consideration each of the above aspects.

7 Conclusion

Sensing mechanisms have evolved over the years to provide accurate monitoring and actuation capabilities within a building. However, it requires an expert user to connect different sensing technologies together to develop applications on top of it. In order to enable easier development of building level applications, a building management system encompassing wireless sensor network services is required. HomeOS, Sensor Andrew and sMAP are some of the few attempts towards this direction. Deployment of wireless sensor network within buildings provide more challenges than already established standards for deploying WiFi and wired devices. Wireless connectivity, aesthetics and remote maintenance are just some of the challenges, and they get exacerbated as the number of devices grow quickly with different applications. We are still a long way away from the *smart buildings* vision of automation within buildings, but we are surely making progress to make it a reality in the future.

References

[1] Heimann sensor. `http://www.heimannsensor.com/`.

[2] Semantic web. `http://www.w3.org/standards/semanticweb`.

[3] Sensource inc. `http://www.sensourceinc.com/`.

[4] Z-Wave.com - ZwaveStart. `http://www.z-wave.com/`.

[5] ZigBee Alliance. `http://www.zigbee.org/`.

[6] AGARWAL, Y., BALAJI, B., DUTTA, S., GUPTA, R., AND WENG, T. Duty-cycling Buildings Aggressively: The Next Frontier in HVAC Control. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on* (2011), IEEE, pp. 246–257.

[7] AGARWAL, Y., BALAJI, B., GUPTA, R., LYLES, J., WEI, M., AND WENG, T. Occupancy-driven Energy Management for Smart Building Automation. In *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building* (2010), ACM, pp. 1–6.

[8] AGARWAL, Y., GUPTA, R., DOMAKI, D., AND WENG, T. Building Depot: An Extensible and Distributed Architecture for Building Data Storage, Access and Sharing. In *ACM Workshop on Embedded Systems For Energy-Efficiency In Buildings(Buildsys)* (2012).

[9] AGARWAL, Y., WENG, T., AND GUPTA, R. The Energy Dashboard: Improving the Visibility of Energy Consumption at a Campus-wide Scale. In *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings* (2009), ACM, pp. 55–60.

[10] Barnich, O., and Van Droogenbroeck, M. ViBe: A Powerful Random Technique to Estimate the Background in Video Sequences. In *Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on* (2009), Ieee, pp. 945–948.

[11] Barrenetxea, G., Ingelrest, F., Schaefer, G., and Vetterli, M. The Hitchhiker's Guide to Successful Wireless Sensor Network Deployments. In *Proceedings of the 6th ACM conference on Embedded network sensor systems* (2008), ACM, pp. 43–56.

[12] Beckel, C., Sadamori, L., and Santini, S. Towards Automatic Classification of Private Households Using Electricity Consumption Data. In *To appear in the 4th ACM Workshop On Embedded Sensing Systems For Energy-Efficiency In Buildings (BuildSys' 12). Toronto, Canada* (2012).

[13] Bushby, S. BACnet™: A Standard Communication Infrastructure for Intelligent Buildings. *Automation in Construction 6*, 5 (1997), 529–540.

[14] Bushby, S., and Newman, H. BACnet Today. *ASHRAE journal 10* (2002), 10–18.

[15] Ceri, S., Gottlob, G., and Tanca, L. What you always wanted to know about Datalog (and never dared to ask). *Knowledge and Data Engineering, IEEE Transactions on 1*, 1 (1989), 146–166.

[16] Dawson-Haggerty, S., Jiang, X., Tolle, G., Ortiz, J., and Culler, D. sMAP: A Simple Measurement and Actuation Profile for Physical Information. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems* (2010), ACM, pp. 197–210.

[17] Dawson-Haggerty, S., Lanzisera, S., Taneja, J., Brown, R., and Culler, D. @scale: Insights from a Large, Long-lived Appliance Energy WSN. In *Proceedings of the 11th international conference on Information Processing in Sensor Networks* (2012), ACM, pp. 37–48.

[18] Dixon, C., Agarwal, R., Saroiu, A., and Bahl, P. An Operating System for the Home. In *Proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation* (2012).

[19] Enck, W., Ongtang, M., and McDaniel, P. Understanding Android Security. *Security & Privacy, IEEE 7*, 1 (2009), 50–57.

[20] Erickson, V., Carreira-Perpiñán, M., and Cerpa, A. OBSERVE: Occupancy-based System for Efficient Reduction of HVAC Energy. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on* (2011), IEEE, pp. 258–269.

[21] Evans, R., Blotter, J., and Stephens, A. Flow Rate Measurements using Flow-induced Pipe Vibration. *Journal of fluids engineering 126*, INEEL/JOU-03-00625 (2004).

[22] Falconer, D., Ficklin, R., and Konolige, K. Robot-mounted Through-wall Radar for Detecting, Locating, and Identifying Building Occupants. In *Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on* (2000), vol. 2, IEEE, pp. 1868–1875.

[23] Farinaccio, L., and Zmeureanu, R. Using a Pattern Recognition Approach to Disaggregate the Total Electricity Consumption in a House into the Major End-uses. *Energy and Buildings 30*, 3 (1999), 245–259.

[24] Fogarty, J., Au, C., and Hudson, S. Sensing from the Basement: A Feasibility Study of Unobtrusive and Low-cost Home Activity Recognition. In *Proceedings of the 19th annual ACM symposium on User interface software and technology* (2006), ACM, pp. 91–100.

[25] Froehlich, J., Larson, E., Campbell, T., Haggerty, C., Fogarty, J., and Patel, S. HydroSense: Infrastructure-mediated Single-point Sensing of Whole Home Water Activity. In *Proceedings of the 11th international conference on Ubiquitous computing* (2009), ACM, pp. 235–244.

[26] Fu, Z., and Culurciello, E. A 1.2 mW CMOS Temporal-difference Image Sensor for Sensor Networks. In *Circuits and Systems, 2008. ISCAS 2008. IEEE International Symposium on* (2008), IEEE, pp. 1064–1067.

[27] Hart, G. Non-intrusive Appliance Load Monitoring. *Proceedings of the IEEE 80*, 12 (1992), 1870–1891.

[28] Hnat, T., Srinivasan, V., Lu, J., Sookoor, T., Dawson, R., Stankovic, J., and Whitehouse, K. The Hitchhiker's Guide to Successful Residential Sensing Deployments. In *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems* (2011), ACM, pp. 232–245.

[29] Jiang, X., Dawson-Haggerty, S., Dutta, P., and Culler, D. Design and Implementation of a High-fidelity AC Metering Network. In *International Conference on Information Processing in Sensor Networks (IPSN).* (2009), IEEE, pp. 253–264.

[30] Kamthe, A., Jiang, L., Dudys, M., and Cerpa, A. SCOPES: Smart Cameras Object Position Estimation System. *Wireless Sensor Networks* (2009), 279–295.

[31] Kazandjieva, M., Gnawali, O., Heller, B., Levis, P., and Kozyrakis, C. Identifying Energy Waste Through Dense Power Sensing and Utilization Monitoring. Tech. rep., Stanford University, Tech. Rep. CSTR 2010-03, 2010.

[32] Kim, Y., Schmid, T., Charbiwala, Z., Friedman, J., and Srivastava, M. NAWMS: Nonintrusive Autonomous Water Monitoring System. In *Proceedings of the 6th ACM conference on Embedded network sensor systems* (2008), ACM, pp. 309–322.

[33] Kim, Y., Schmid, T., Charbiwala, Z., and Srivastava, M. ViridiScope: Design and Implementation of a Fine-grained Power Monitoring System for Homes. In *Proceedings of the 11th international conference on Ubiquitous computing* (2009), ACM, pp. 245–254.

[34] Krioukov, A., Dawson-Haggerty, S., Lee, L., Rehmane, O., and Culler, D. A Living Laboratory Study in Personalized Automated Lighting Controls. In *ACM Workshop on Embedded Systems For Energy-Efficiency In Buildings(Buildsys)* (2011).

[35] Langendoen, K., Baggio, A., and Visser, O. Murphy Loves Potatoes: Experiences from a Pilot Sensor Network Deployment in Precision Agriculture. In *Parallel and Distributed Processing Symposium, 2006. IPDPS 2006. 20th International* (2006), IEEE, pp. 8–pp.

[36] Laughman, C., Lee, K., Cox, R., Shaw, S., Leeb, S., Norford, L., and Armstrong, P. Power Signature Analysis. *Power and Energy Magazine, IEEE 1*, 2 (2003), 56–63.

[37] Leeb, S., Shaw, S., and Kirtley Jr, J. Transient Event Detection in Spectral Envelope Estimates for Nonintrusive Load Monitoring. *Power Delivery, IEEE Transactions on 10*, 3 (1995), 1200–1210.

[38] Li, L., Huang, W., Gu, I., and Tian, Q. Foreground Object Detection from Videos Containing Complex Background. In *Proceedings of the eleventh ACM international conference on Multimedia* (2003), ACM, pp. 2–10.

[39] Lifton, J., Feldmeier, M., Ono, Y., Lewis, C., and Paradiso, J. A Platform for Ubiquitous Sensor Deployment in Occupational and Domestic Environments. In *Information Processing in Sensor Networks, 2007. IPSN 2007. 6th International Symposium on* (2007), IEEE, pp. 119–127.

[40] Marchiori, A., and Han, Q. Distributed Wireless Control for Building Energy Management. In *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building* (2010), ACM, pp. 37–42.

[41] Martocci, J. BACnet® Unplugged: ZigBee® and BACnet Connect. *ASHRAE Journal 50*, 6 (2008), 42–47.

[42] Murakita, T., Ikeda, T., and Ishiguro, H. Human Tracking using Floor Sensors Based on the Markov Chain Monte Carlo Method. In *Pattern Recognition, 2004. ICPR 2004. Proceedings of the 17th International Conference on* (2004), vol. 4, IEEE, pp. 917–920.

[43] Orr, R., and Abowd, G. The Smart Floor: A Mechanism for Natural User Identification and Tracking. In *CHI'00 extended abstracts on Human factors in computing systems* (2000), ACM, pp. 275–276.

[44] Park, T., Chon, Y., Park, D., and Hong, S. BACnet over ZigBee, A New Approach to Wireless Datalink Channel for BACnet. In *Industrial Informatics, 2007 5th IEEE International Conference on* (2007), vol. 1, IEEE, pp. 33–38.

[45] Patel, S., Robertson, T., Kientz, J., Reynolds, M., and Abowd, G. At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line. In *Proceedings of the 9th international conference on Ubiquitous computing* (2007), Springer-Verlag, pp. 271–288.

[46] Penn, D., Oberzaucher, E., Grammer, K., Fischer, G., Soini, H., Wiesler, D., Novotny, M., Dixon, S., Xu, Y., and Brereton, R. Individual and Gender Fingerprints in Human Body Odour. *Journal of the Royal Society Interface 4*, 13 (2007), 331–340.

[47] Perez-Lombard, L., Ortiz, J., and Pout, C. A Review on Buildings Energy Consumption Information. *Energy and Buildings 40*, 3 (2008), 394–398.

[48] Potamitis, I., Chen, H., and Tremoulis, G. Tracking of Multiple Moving Speakers with Multiple Microphone Arrays. *Speech and Audio Processing, IEEE Transactions on 12*, 5 (2004), 520–529.

[49] Prudenzi, A. A Neuron Nets Based Procedure for Identifying Domestic Appliances Pattern-of-use from Energy Recordings at Meter Panel. In *Power Engineering Society Winter Meeting, 2002. IEEE* (2002), vol. 2, IEEE, pp. 941–946.

[50] Roth, K., Dieckmann, J., and Brodrick, J. Demand Control Ventilation. *ASHRAE journal 45*, 7 (2003), 91–92.

[51] Rowe, A., Berges, M., Bhatia, G., Goldman, E., Rajkumar, R., Garrett, J., Moura, J., and Soibelman, L. Sensor Andrew: Large-scale Campus-wide Sensing and Actuation. *IBM Journal of Research and Development 55*, 1.2 (2011), 6–1.

[52] Shelby, Z., and Bormann, C. *6LoWPAN: The Wireless Embedded Internet*, vol. 43. Wiley, 2011.

[53] Simmhan, Y., Prasanna, V., Aman, S., Natarajan, S., Yin, W., and Zhou, Q. Towards Data-driven Demand Response Optimization in a Campus Microgrid. In *ACM Workshop On Embedded Sensing Systems For Energy-Efficiency In Buildings. ACM* (2011).

[54] Szewczyk, R., Polastre, J., Mainwaring, A., and Culler, D. Lessons from a Sensor Network Expedition. *Wireless Sensor Networks* (2004), 307–322.

[55] Teixeira, T., Dublon, G., and Savvides, A. A Survey of Human-Sensing: Methods for Detecting Presence, Count, Location, Track, and Identity. *ACM Computing Surveys 5* (2010).

[56] Toda, H., and Capi, G. Development of a High-Resolution Human-Specific Breath Gas Sensor for Survivor Detection in Disaster Zones. *Advanced Robotics 26*, 3-4 (2012), 349–362.

[57] US Department of Energy. Buildings Energy Data Book. `http://buildingsdatabook.eren.doe.gov/`, 2010.

[58] Valtonen, M., Maentausta, J., and Vanhala, J. TileTrack: Capacitive Human Tracking using Floor Tiles. In *Pervasive Computing and Communications, 2009. PerCom 2009. IEEE International Conference on* (2009), IEEE, pp. 1–10.

[59] Weng, T., Balaji, B., Dutta, S., Gupta, R., and Agarwal, Y. Managing Plug-Loads for Demand Response within Buildings. In *ACM Workshop on Embedded Systems For Energy-Efficiency In Buildings(Buildsys)* (2011).

[60] Werner-Allen, G., Lorincz, K., Ruiz, M., Marcillo, O., Johnson, J., Lees, J., and Welsh, M. Deploying a Wireless Sensor Network on an Active Volcano. *Internet Computing, IEEE 10*, 2 (2006), 18–25.

[61] Workgroup, U. E. P. A. G. B. Buildings and their Impact on the Environment: A Statistical Summary. In *http://buildingsdatabook.eren.doe.gov/* (2009), Department of Energy (DOE).

[62] Zetik, R., Crabbe, S., Krajnak, J., Peyerl, P., Sachs, J., and Thoma, R. Detection and Localization of Persons Behind Obstacles using M-sequence Through-the-wall Radar. In *Proc. SPIE* (2006), vol. 6201, Citeseer.

[63] Zhou, Q., Liu, J., Host-Madsen, A., Boric-Lubecke, O., and Lubecke, V. Detection of Multiple Heartbeats using Doppler Radar. In *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on* (2006), vol. 2, IEEE, pp. II–II.

[64] Zhou, Q., Natarajan, S., Simmhan, Y., and Prasanna, V. Semantic Information Modeling for Emerging Applications in Smart Grid. In *Ninth International Conference on Information Technology: New Generations (ITNG)* (2012), IEEE.