ePrints: A Real-Time and Scalable System for Fair Apportionment and Tracking of Personal Energy Footprints in Commercial Buildings

Peter Wei
Columbia University

Xiaoqi Chen
Tsinghua University

Jordan Vega
Columbia University

Stephen Xia
Columbia University

Rishikanth Chandrasekaran
Columbia University

Xiaofan Jiang
Columbia University

ABSTRACT

We propose a system that tracks each occupant's personal share of energy use, or "energy footprint", inside commercial building environments, and provides insights to occupants on the real-time energy impact of their actions. We propose a new space-centric policy for fair apportionment of energy in shared environments and demonstrate a method for automatically determining space-centric energy zones. We design and implement ePrints – a system for tracking personalized energy usage in real-time. ePrints supports different apportionment policies, with μs-level footprint computation time and graceful scaling with size of building, frequency of energy updates, and rate of occupant location changes. Finally, we present applications enabled by our system, such as mobile and wearable applications to provide users timely feedback on the energy impacts of their actions, as well as applications to provide energy saving suggestions and inform building-level policies.

CCS CONCEPTS

•**Human-centered computing** →**Ubiquitous and mobile computing;** •**Computer systems organization** →*Real-time systems;* •**Spatial-temporal systems** →Location based services;

KEYWORDS

Energy Apportionment, Commercial Buildings, Energy Footprint, Energy Zoning, System Scalability

ACM Reference format:

Peter Wei, Xiaoqi Chen, Jordan Vega, Stephen Xia, Rishikanth Chandrasekaran, and Xiaofan Jiang. 2017. ePrints: A Real-Time and Scalable System for Fair Apportionment and Tracking of Personal Energy Footprints in Commercial Buildings. In *Proceedings of The 4th International Conference on Systems for Energy-Efficient Built Environments, Delft, The Netherlands, November 8-9, 2017 (BuildSys),* 10 pages.
DOI: 10.475/123_4

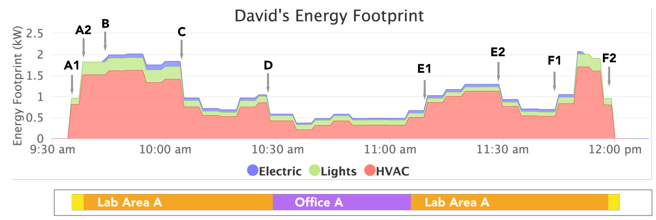

Figure 1: Real-time energy footprint of an occupant. The colored stripes indicate different spaces where the occupant is located (yellow refers to the hallway).

1 INTRODUCTION

Buildings consume a large portion of the total electricity in the United States. Products such as Nest smart thermostats have emerged in recent years to help people reduce their daily energy consumption in homes. While the home is a sensible area to focus, we spend the majority of our active moments during the day inside commercial buildings. Consequently, the greatest opportunity for us to reduce energy usage is when we are at work and inside our offices.

However, promoting energy savings in commercial and office buildings comes with a different set of challenges than in a residential setting. Unlike in residential areas where families pay their own energy bills, there is little incentive for occupants to save energy in commercial buildings, as the energy bills are generally paid by employers and companies. Also, while commercial buildings are becoming smarter with increasing numbers of energy monitoring endpoints from advances in building energy monitoring, the effect of an occupant's personal actions on the overall energy consumption of the building remains unknown. This is because commercial building rooms, spaces, and appliances are generally shared among multiple occupants, so the non-trivial problem of fairly apportioning and attributing the correct energy consumption to each occupant is exacerbated compared to residential homes. As a result, an individual, who is ignorant of his or her own consumption within the building, is less motivated to save energy. In this paper, we propose a system that tracks each occupant's unique and individualized energy usage, or **energy footprint**, in real-time in commercial building environments. Our system provides insights to occupants on the *real-time energy impact of their actions.*

Figure 1 shows the annotated footprint of one occupant in our system over the course of a morning: (A1) David enters the deployment area, and ePrints associates the energy consumption of the hallway environment (primarily HVAC and lighting) with him and calculates his energy footprint. (A2) He moves to lab area A, and his footprint changes to reflect his consumption in the new area. (B) David powers on personal energy resources. (C) Other occupants enter the lab area, thus reducing David's energy footprint.

DOI: 10.475/123_4

(D) David moves to Office A for a meeting, and shares the energy consumption with the other occupants in the office. (E1) After the meeting, he returns to lab A; his energy footprint increases as a response. (E2) David receives a suggestion to reduce his energy consumption by changing the temperature setpoint of the room. (F1) Occupants in lab A begin to leave for lunch; David, who is staying later, begins to see an increase in his energy footprint as a larger share of HVAC and lighting is apportioned to him. (F2) Finally, David leaves the lab and his footprint drops to zero.

In this work, we present the following contributions:

- We present the design and implementation of a **scalable** energy footprinting system that provides **real-time energy footprints** to occupants in shared environments. Algorithms that propagate energy footprint changes, combined with a tripartite graph that utilizes spaces as an intermediary stage between energy resources and occupants, enables a low-latency system that scales linearly with occupants and energy monitoring updates.
- To **fairly** apportion energy inside shared environments, we proposes a new apportionment strategy based on the notion of "spaces" where spaces are defined based on **human-centric zones** instead of traditional HVAC zones.
- We further introduce a **policy manager** that supports diverse apportionment strategies, including space-level and individual-level apportionment policies. This policy manager is critical for ePrints to be adopted in a variety of shared environments.
- We deploy ePrints in a shared environment, perform studies with occupants, collect energy footprint data, and evaluate the system on the fairness of multiple policies to demonstrate the advantages of ePrints in a shared environment.
- We demonstrate ePrints' potential to enable other applications, such as energy saving suggestions and building-level policies.

We leverage and build upon existing building energy work including building energy monitoring and indoor localization. We do **not** claim novelty or contributions in these two areas, instead focusing on the energy apportionment policy, architecture, algorithm, and system design to fairly and efficiently apportion energy inside large buildings into personal energy footprints. We further demonstrate the potential of ePrints as a platform to enable diverse human-centric building energy applications.

2 RELATED WORK

This project draws from and builds upon a diverse set of research areas, including occupancy detection, indoor localization, building energy monitoring, and energy apportionment. However, the focus of this project is not to make improvements to indoor localization or energy monitoring, but to utilize them as building blocks as part of a larger system that fairly and efficiently apportions energy inside commercial buildings into individualized *energy footprints*.

Energy monitoring in commercial buildings often includes monitoring of miscellaneous electric loads (MELs), lighting, and HVAC (heating, ventilation and air conditioning). In MELs monitoring, plug-load meters, both wired [9] and wireless [11, 17], have been used to monitor plug-loads directly. Plug-meters are accurate and simple to deploy in homes. However, they become less practical in large buildings with thousands of electric outlets. An intermediate solution is to monitor electricity usage at the circuit level, by tapping onto circuit breakers [22]. This approach is more cost-effective than plug-meters, but provides a less granular view of energy usage. At the other end of the spectrum, non-intrusive load monitoring (NILM)-based approaches use algorithmic techniques to estimate plug-level usage from a single high-fidelity power meter at the house or building level [14], some require training [20] or additional sensor inputs [15]. While the cost of NILM-based approaches is lower, accuracy is often less than desirable. Lighting and HVAC in commercial buildings can be monitored directly by connecting to the building management system (BMS) through protocols such as BACNet, LonTalk, and Modbus [3]. Recent works have also demonstrated that equivalent energy consumption of HVAC resources can be estimated using data feeds, such as fan speeds and valve positions, combined with physical HVAC zone models [2]. However, BMS is often unavailable or inaccessible without going through layers of administrative approval. To obtain a holistic view of building energy use, ePrints interfaces with a mixture of underlying building energy monitoring technologies to provide full coverage of energy resources.

Occupant localization is another key component of ePrints. Various indoor localization technologies have been proposed over the past two decades. Due to the ubiquity of smartphones, many popular solutions are based on received signal strengths (e.g. RSSI) of WiFi/Bluetooth signals [1, 13, 19]. In general, WiFi-based systems can achieve good accuracy and can utilize existing infrastructure in commercial buildings [27]; however, placement of WiFi access points (APs) is typically optimized for wireless coverage, not localization. In comparison, Bluetooth Low Energy (BLE) beacons are inexpensive, battery-powered, and can be optimally deployed to maximize localization accuracy. Previous works have also shown that BLE methods outperform WiFi-based methods in terms of indoor localization accuracy [27]. Many other types of indoor localization technologies, including vision-based [7] and acoustic-based methods [10, 24], have been proposed with varying degrees of accuracy, cost, and ease of deployment. Magnetic field strengths and modulated magneto-inductive beacons can also be used to estimate object locations [12, 18, 23] with high accuracy and precision. ePrints works with a diverse set of localization technologies that provide indoor "space-level" occupant location accuracy.

Apportionment strategy is the policy that determines how energy is divided among occupants in shared environments and can have drastic effects on the fairness and incentives to push occupants and building managers to reduce energy usage. Since there is often no single correct solution for apportionment, research in this area has provided numerous methods for dividing energy consumption among occupants within a physical space [8, 26]. Apportionment strategies in previous works typically apportion energy consumption of a physical HVAC/lighting zone directly to its occupants. However, in large shared environments, an HVAC/lighting zone is often divided into sub-spaces occupied by distinct groups of people, where each sub-space has a fixed area/volume and its own energy-use preference or requirement (e.g. different departments sharing a common cubicle space; two research groups with different temperature preferences sharing a single lab area).

In recent years, a few research projects have combined energy monitoring, occupant localization, and energy apportionment to estimate personal energy consumption in real-time. These systems focus on residential homes or buildings where rooms are owned individually. For example, [16] and [25] monitor energy consumption in residential homes and in a dormitory setting to understand the energy consumption of each occupant. However, in large commercial buildings such as office spaces, where one room or space could house hundreds of employees belonging to different departments or having different functions, a different monitoring system is required to accurately monitor and fairly apportion energy consumption. We address these challenges by introducing the notion of human-centric zoning, described in more detail in section 4.1.1.

3 CHALLENGES

To ensure a fair and timely energy footprint in commercial buildings, two challenges must be addressed: how to fairly apportion energy consumption to individuals, and how to design scalable and efficient algorithms to enable energy footprints in real-time.

3.1 Energy Apportionment Policy

Hay et al. [8] defines energy apportionment as "a process of dividing up the total consumption of a building... and allocating it to individuals". In commercial buildings, the total consumption can be thought of as the aggregate of all energy consuming resources, including HVAC, lighting, and plug loads. A personal energy footprint will display the total energy consumption that a single occupant is apportioned at a given time. To illustrate the challenges in defining fair energy apportionment policies, we present two realistic scenarios in commercial buildings.

Our first example is a single space that is jointly shared by two companies and is cooled by a single HVAC unit. Company A occupies 20% of the space and has two employees, while company B occupies the other 80% of the space and has two employees. Common apportionment policies such as dividing the HVAC consumption evenly over the occupants may succeed at the individual level, but fail to incorporate the uneven split of the company areas. Although the HVAC consumption is allocated equally among the occupants (25%), there is an unfair split of consumption due to the uneven split in space that the employees occupy.

The second example is the case of office worker A in his personal cubicle that comes equipped with his personal desktop computer. Worker A is the only person who uses this desktop. If worker B comes to visit worker A at his cubicle, some of worker A's desktop consumption may be allocated to worker B, unless there is a policy that associates the device to worker A. As a result, worker B will be unfairly apportioned a portion of worker A's desktop consumption.

The two examples presented show that apportionment fairness varies depending on space organization and function. First, commercial buildings often have rooms, open areas, and resources that are shared among different groups or organizations; as such, dividing spaces based on the physical properties of the building (e.g. rooms in the building floor plan or HVAC zones), as is often done in residential areas, is not sufficient to ensure fair apportionment. Second, different spaces have different measures of fairness depending on the function of the space (e.g. personal office space vs. a public area). It is thus imperative to design a system that is flexible to different energy apportionment policies and adaptable to future changes in apportionment policy.

3.2 Real-Time and Scalable Energy Footprint Computation

The second challenge arises from the need of a scalable system that can provide real-time energy footprint updates. Personal footprints updates should be *available* to the user within a few seconds, at most, to ensure an accurate footprint; it is important to note that the delay between when an energy change occurs and when this change becomes *visible* to the user depends on how often the user client (e.g. a smartphone) polls for updates. Additionally, ePrints must meet the timing requirements even with a large number of devices, spaces, and occupants. The rest of the section introduces the metrics we use to evaluate latency and scalability.

In our system, there are two events that cause changes in personal energy footprints: energy resource consumption changes and occupant location changes. We refer to the latency of each event as energy change latency and location change latency, respectively, and define each latency as the time difference between the change and the feedback to the occupant. The energy change latency is the sum of three sources of latency: energy monitoring response time, energy footprint computation, and feedback latency. The energy monitoring response time is dependent on the rate that energy monitoring nodes report updates and is discussed in Section 5.2. The feedback latency, or the time between the completion of the footprint computation and the feedback medium (mobile device, web portal, wearable) notification, is discussed in Section 5.4.2.

Location change latency is also a sum of three sources of latency: localization delay, energy footprint computation, and feedback latency. The localization delay, or the delay between the physical change in location of an occupant and the acknowledgement of the server of the location change, is discussed in Section 5.3.

The energy footprint computation latency, or the time to compute an individual's energy footprint, is critical to the scalability of ePrints. Especially in commercial buildings with large numbers of occupants, energy consuming resources, and spaces, the energy footprint computation latency should remain low. The algorithms chosen must also scale gracefully to any building size.

For ePrints to be scalable, personal energy footprinting queries must have low computational complexity, and all computations must be scalable to any building size. The system must be capable of servicing hundreds of footprinting queries per second with low latency even in deployments with thousands of energy consuming resources, hundreds of zones, and thousands of people.

4 SYSTEM DESIGN

In this section, we propose new apportionment and system concepts to address the challenges from Section 3. To address the issues of energy apportionment in commercial buildings, we propose three novel concepts: a new scheme for partitioning building spaces based on space ownership or human behavior, which we term "human-centric zoning"; a tripartite graph representation for storing energy data, as well as the idea of space-level and individual-level policies for dividing consumption of energy resources; and a policy manager to accommodate different types of apportionment policies.

To address the latency and scalability of ePrints, we develop and analyze the time complexity of algorithms for energy updates and queries over the proposed tripartite graph data structure.

4.1 Addressing Energy Apportionment

4.1.1 Human-Centric Zoning. As mentioned in Section 3.1, one of the primary challenges for fair energy apportionment in commercial buildings is the prevalence of rooms, spaces, and resources that are shared between multiple groups. To address this issue, we propose the idea of human-centric zoning, where spaces are partitioned based on either space ownership or occupant behavior rather than physical boundaries, like rooms. We discuss two methods of accomplishing this partitioning.

Human-centric zoning based on space ownership can be accomplished by dividing rooms and areas based on the spaces assigned to each group or entity. In the case of the first example mentioned in Section 3.1, rather than denoting the entire room as a single space, we define two spaces within the room where the employees of company A and company B reside. As a result, the occupants are only responsible for consumption within the portion of the room that belongs to their company. However, a priori knowledge of space ownership is required to partition using this method.

If a priori knowledge of space ownership is not known, then human-centric zoning can still be accomplished by partitioning based on mobility patterns (location traces) of occupants. This requires occupants to be localized using indoor localization technologies. For ePrints, we used BLE beacons due to their ease of deployment and low-cost. Additionally, their room/cubicle-level accuracy is sufficient for our application. We distributed BLE beacons throughout our deployment space to obtain complete physical coverage. Bluetooth fingerprints (e.g. RSSI values) are collected and reported to the server by the *Energy Footprint* application installed on the occupant's smartphone. This allows us to collect occupant location traces in coordinate space (x, y), following standard fingerprint localization algorithms such as in [21]. Once a significant amount of coordinates are collected, one can cluster the coordinate data to determine partitions of the space using any clustering algorithm and any method for determining the number of clusters.

4.1.2 Space-Level and Individual-Level Policies. Once building spaces are partitioned via human-centric zoning, we use the tripartite graph representation, $K_{D,S,P}$, shown in Figure 2, to organize the sets of energy resources (D), spaces (S), and occupants (P). In the tripartite graph, an edge exists between zone $s \in S$ and resource $d \in D$ if the resource influences the zone, and an edge exists between occupant $p \in P$ and zone $s_i \in S$ if the occupant is localized to the zone. The neighborhood $\mathcal{N}$ of a node is thus the set of all other nodes connected via edge to that node. The primary benefit of the tripartite graph model is the complete dissociation of shared resources from the occupants. This dissociation enables us to propose the idea of **space-level** and **individual-level** policies for fair apportionment of shared resources in commercial buildings.

Recall that in commercial buildings, the issue of fair energy apportionment arises because different groups of people are often within the same physical boundary that common apportionment schemes use to determine spaces. Using human-centric zoning, the building is partitioned based on space ownership of different groups to resolve this issue. However, there are now energy resources that are shared between different spaces. To resolve this new issue, we propose to first apportion resources over each space

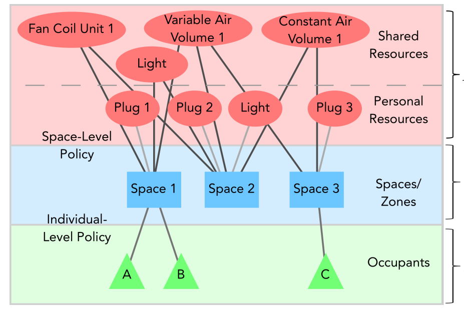

Figure 2: Illustration of the tripartite graph $K_{D,S,P}$.

using space-level policies. Then, using individual-level policies, energy apportioned to a space is further divided among the occupants. If we define the space apportionment policy function, $f(s, d)$, and the individual policy function, $g(p, s, d)$, we can define the total apportionment policy function (Equation 1), which refers to the aggregate apportionment policy.

$$h(p, s, d) = f(s, d)g(p, s, d) \tag{1}$$

Splitting the policy into space-level apportionment and individual-level apportionment allows for more flexibility and more diverse apportionment schemes for different situations. Our zoning model and the concept of dual-level policies enables our footprinting system to achieve fairness among spaces/groups as well as fairness among individuals using shared energy resources.

4.1.3 Policy Manager. Previously, we presented the idea of dual-level apportionment policies that enable us to independently implement fair policies over groups of people, common spaces, and occupants within each space. However, different subsets of spaces, devices, or individuals may require a different set of policies to satisfy fairness; to address this issue, we introduce the idea of a policy manager that allows managers to implement different space-level and individual-level policies for different spaces, devices, and individuals depending on the situation and fairness requirements.

4.2 Scalable Energy Footprinting Algorithms

A system capable of computing energy footprints in real-time requires algorithms that are both efficient as well as scalable. As touched upon in [26], simpler apportionment policies are not as desirable as more complex policies, because they cannot fully capture the specific interactions between people and the building; however, simpler apportionment policies tend to surpass more complex policies in terms of runtime, which can make a significant difference if the policy is recomputed many times.

In our system, there are three common types of possible operations performed on the tripartite graph: energy resource consumption change, occupant location change, and occupant footprint calculation. When a device's energy consumption changes or an occupant changes location, we are faced with a decision: to update the affected occupants' energy footprints immediately, or postpone the update until the occupant requests their footprint. We designed two sets of algorithms for the updates and footprint queries, one that immediately propagates the change to the occupants (called **propagation**), and one that postpones the footprint calculation until the occupant footprint query occurs (called **delayed update**).

Algorithm 1: Occupant Location Change with Propagation

procedure LOCATIONCHANGEUPDATE($P, s_{source}, s_{dest}, D, p$)

$D_{source} \leftarrow \mathcal{N}(s_{source}) \cap D$ ▷ devices influencing the space

$P_{source} \leftarrow \mathcal{N}(s_{source}) \cap P \setminus p$ ▷ occupants in the space except person changing location

for $p_i \in P_{source}$ **do**

$\pi_{p_i} \leftarrow 0$ ▷ π_{p_i} is the energy apportioned to person p_i

for $d \in D_{source}$ **do**

$UpdatePolicy(g, p_i, s_{source}, d)$

$\pi_{p_i} \leftarrow \pi_{p_i} + f(s_{source}, d)g(p_i, s_{source}, d)E(d)$

$D_{dest} \leftarrow \mathcal{N}(s_{dest}) \cap D$

$P_{dest} \leftarrow \mathcal{N}(s_{dest}) \cap P \setminus p$

for $p_j \in P_{dest}$ **do**

$\pi_{p_j} \leftarrow 0$

for $d \in D_{dest}$ **do**

$UpdatePolicy(g, p_j, s_{dest}, d)$

$\pi_{p_j} \leftarrow \pi_{p_j} + f(s_{dest}, d)g(p_j, s_{dest}, d)E(d)$

Algorithm 2: Resource Energy Change with Propagation

procedure ENERGYCONSUMINGRESOURCEUPDATE($d, S, P, \Delta E$)

for $s \in \mathcal{N}(d)$ **do** ▷ Spaces influenced by device

for $p \in \mathcal{N}(s) \cap P$ **do** ▷ People localized to space

$UpdatePolicy(f, s, d)$

$\pi_p \leftarrow \pi_p + f(s, d)g(p, s, d)\Delta E$

Algorithm 3: Energy Footprint Query (delayed update)

procedure FOOTPRINTQUERY(p)

$s \leftarrow Space(p)$

$\pi_p \leftarrow 0$

for $d \in \mathcal{N}(s)$ **do**

$UpdatePolicy(f, s, d)$

$UpdatePolicy(g, p, s, d)$

$\pi_p \leftarrow \pi_p + f(s, d)g(p, s, d)E(d)$

Table 1 presents the runtime of each algorithm (the algorithms with runtime greater than constant time are shown in Algorithms 2, 1, and 3). We define the branching factor (b) as the average number of energy consuming resources associated with each space.

Algorithm	Runtime Complexity	Runtime in our System
Resource Update with Propagation	$O(\frac{b\|P\|}{\|D\|})$	$1.446\mu s$
Location Change with Propagation	$O(\frac{b\|P\|}{\|S\|})$	$6.735\mu s$
Footprint Computation (Propagation)	$O(1)$	**$422ns$**
Resource Update with Delayed Update	$O(1)$	$432ns$
Location Change with Delayed Update	$O(1)$	$1.418\mu s$
Footprint Computation (Delayed Update)	$O(b)$	**$2.254\mu s$**

Table 1: Runtime complexity and system runtime of the algorithms in ePrints.

In our testbed, we assume that the number of times an occupant changes location or a resource consumption changes is much less frequent than energy footprint queries; thus, we implement algorithms with propagation for reduced footprint computation latency. This is validated in Section 6.4, where we evaluate and compare the different algorithms in realistic deployment scenarios.

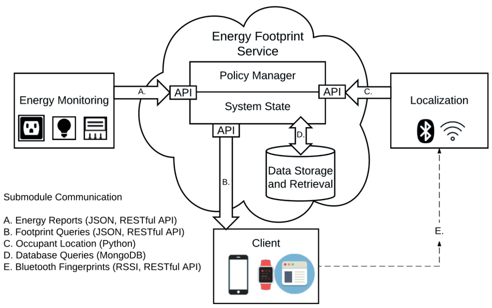

Figure 3: System architecture block diagram.

5 ARCHITECTURE AND IMPLEMENTATION

As shown in Figure 3, ePrints is composed of four loosely-coupled subsystems that work together to provide real-time, personal energy footprints. The *Energy Footprinting Service*, which includes the *Policy Manager*, *System States* and *Data Storage and Retrieval*, reside in the cloud, and the components communicate with each other via standardized interfaces. The Energy Footprinting Service exposes standardized interfaces and APIs that communicate with clients such as Android and iOS applications installed on occupants' mobile devices, as well as a variety of energy monitoring devices including wireless power meters, indirect sensing nodes and BACnet monitors placed around the deployment area.

5.1 Energy Footprinting Service

The Energy Footprinting Service is the central subsystem responsible for receiving localization and energy monitoring data, computing personal energy footprinting data according to the apportionment policies, and exposing the available data to client-side applications. This service is composed of the policy manager responsible for apportionment policies, and the state manager responsible for energy footprinting computations and data storage.

5.1.1 Policy Manager. The policy manager is central to the fair apportionment of energy to the occupants. It is able to accommodate different types of apportionment policies, as discussed in Section 4.1. The policy manager currently manages 4 policies: two space level policies – equal division (equation 2) and proportional to volume (equation 3), as well as two individual level policies – equal division (equation 4) and proportional to isolation (equation 5) [26]. However, the policy manager is flexible such that other policies, such as those discussed in [26], can be easily interchanged.

$$f(s, d) = \frac{1}{|\{s \in \mathcal{N}(d)\}|} \tag{2}$$

$$f(s, d) = \frac{V(s)}{\sum_{s_i \in \mathcal{N}(d)} V(s_i)} \tag{3}$$

$$g(p, s, d) = \frac{1}{|\{p \in \mathcal{N}(s) \cap P\}|} \tag{4}$$

$$g(p, s, d) = \frac{E(p)}{\sum_{p_i \in \mathcal{N}(\mathcal{N}(d)) \cap P} E(p_i)} \tag{5}$$

Here, $V(s)$ refers to the volume of space s, $\mathcal{N}(d)$ refers to the spaces influenced by the device d, and $\mathcal{N}(s) \cap P$ is the set of occupants localized to the space s. The policy manager allows ePrints to be adaptable to a variety of shared environments.

5.1.2 State Management. The system maintains an in-memory data structure containing all energy consuming resources, which include HVAC systems, lighting, and power outlets, as well as occupant statuses and energy zones. It keeps track of the latest states of all energy consumption and the occupants' associations with them. The relationships between energy consuming resources and spaces are initialized using pre-configured metadata, some of which are available in the Building Management System (BMS) while others are not and require manual setup. When an energy consuming resource reports a new power consumption value, or an occupant's location is updated, the in-memory graph changes to reflect the updated energy footprint for all occupants as described by Algorithms 2 and 1 in Section 4.2.

5.1.3 Data Storage and Retrieval. The system periodically stores the latest energy state, which we term a "snapshot", into the database with a corresponding timestamp. While a user can check his most recent energy consumption in the latest state, the historical snapshot can reveal a time-series of the energy consumption, or the more complete energy footprint.

A trade-off exists between higher snapshot sampling frequency and lower storage demand; saving data more frequently can help reconstruction of a more accurate historical footprint, but at the cost of more storage space. We mitigate this dilemma by using a **variable-sampling-rate** scheme. The system stores the snapshot to the database at regular intervals when idle, but stores more frequently when a user's energy footprint changes abruptly, such as during a location change or an energy consuming resource change.

Our implementation uses MongoDB to store the state snapshots, as a NoSQL database is well suited for storing our state data structures. The state is split into three sets, each stored into one collection. We also add the timestamp as an indexing key for the database, as it is a frequent criteria for many queries. The central system also attempts to reconstruct the in-memory state by reading previous snapshots from the database during any failure recovery process.

5.2 Energy Monitoring

Energy monitoring has been studied extensively in the context of commercial buildings. For ePrints, we focus on the following three types of energy monitoring: HVAC monitoring, light sensing, and plug-load monitoring. It is possible to further subdivide these types, but for the purposes of our system, these three types provide sufficient granularity for apportioning energy consumption.

(1) Plug-meters are deployed across the building testbed, including student workstations, offices, and labs. All existing electrical appliances are transferred to metered power strips. They are configured to report at the maximum fixed reporting frequency.
(2) Our system reads building device statuses via BACnet, a building automation protocol. We collaborate with the campus facilities team to obtain a mapping from BACnet IDs to the floor plan in BMS, as well as details of energy-consuming components in the HVAC system. We pull data from these components periodically and report the estimated power consumption.
(3) Some HVAC units are not exposed under BACnet, and the occupancy-aware smart lighting system in our building is not connected to BMS; we monitor these resources with two types of customized indirect sensing nodes. The first, for detecting lighting, consists of a Huzzah Feather board and a TSL 2561 luminosity sensor. The second, for sensing HVAC, consists of an Intel Edison board and a Wind Sensor from Modern Devices. They are configured to report immediately when a change is detected, or at a minimum frequency when idle.

We have adapted the process of determining HVAC energy consumption from [3] to accommodate our building's sensor data, such as the heat transfer equation to determine heating and cooling loss, applying the equation to both fan coil units (FCU) and variable air volume (VAV) terminal units. In addition, we determine the electric consumption due to the air handling units (AHU) by dividing the VAV unit's air flow by the maximum air flow from the AHUs, and multiplying by the maximum power consumption of the AHU.

The energy monitoring subsystem handles incoming energy consumption reports of various formats, each requiring specific adaptations. To limit adaptation effort, the subsystem reports to the cloud service using standardized protocols and data formats (HTTP, RESTful API, and JSON).

Some devices (e.g. commercial plug-meters) can only submit energy updates at a fixed interval, while others (e.g. customized sensing nodes) are more configurable. To ensure real-timeliness, we use an interrupt-driven energy consumption submission whenever possible; that is, an energy monitoring device submits an immediate update upon sensing a large, abrupt change in the power of the energy consuming unit. This allows the energy snapshot to closely follow the significant changes in energy consumption.

5.3 Localization

The localization subsystem can work with different localization technologies, providing a standardized interface to various localization implementations. Due to the ubiquity of smartphones, WiFi fingerprinting and Bluetooth beaconing are two popular indoor localization techniques. We chose Bluetooth due to its ease of deployment, low cost, and ability to meet ePrints' relatively course-grained location accuracy requirement. We deployed 42 BluVision iBeek beacons in the testbed, and trained multiple classifiers to achieve high accuracy and low testing computation time.

The location change latency mentioned in 3.2 is divided into device polling time and server classification time. Each mobile device polls for Bluetooth beacon signal strength values at a set interval. Server classification time of location is dependent on the classifier; although k-Nearest Neighbors (kNN) is the simplest to deploy, we instead implemented Support Vector Machines (SVM) for higher accuracy and lower classification time.

The placement of beacons to maximize localization accuracy while reducing cost is an interesting problem that depends heavily on the environment. We employed a naive strategy for placing beacons; we placed more beacons in locations with dense, smaller sized zones, and less beacons in locations with sparse, larger sized zones. However, prior knowledge of locations of interest allows for more complex strategies such as those described in [4, 5].

5.4 Web, Mobile, and Wearable Applications

Our system supports various client applications via a standard set of APIs. We list some frequently used ones below:

```
historical_footprint(begin_time, end_time, user_id)
```

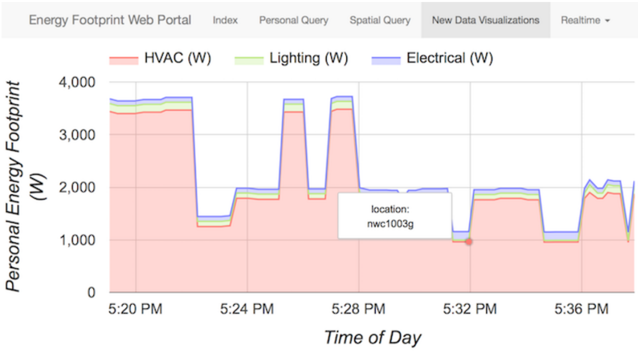

Figure 4: Personal Energy Footprint Web Portal

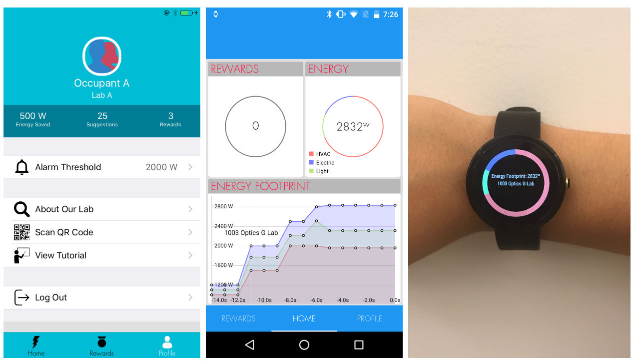

Figure 5: Left to right: iOS application settings tab, Android application energy footprint tab, Android Wear application.

```
    return: [(timestamp, foot_print_val)]
current_power(user_id)
    return: (timestamp, foot_print_val)
submit_location(user_id, localization_data)
    return: user_location
```

Web, mobile and wearable applications demonstrate the potential of ePrints to inform occupants and induce change. We describe them in more detail this section.

5.4.1 Personal Energy Footprint Web Portal. We provide various visualizations of energy footprints through the application's web portal. For example, Figure 4 displays the personal energy footprint of an occupant separated by the resource type. These footprints and visualizations provide the means to analyze and visualize energy consumption behavior of devices, rooms, and occupants.

5.4.2 Mobile. We develop Android and iOS smartphone client application for building occupants to provide timely feedback about the consequences of their actions on their personal energy footprint.

The application collects and submits localization fingerprints to the Energy Footprinting Service at a periodic interval when the mobile device is within the experiment area. The service responds with the occupant's current energy footprint. The feedback latency mentioned in Section 3.2 is dependent on the periodic interval; a longer interval results in longer feedback latency. The client displays the occupant's real-time energy footprint broken down by resource type, as well as a visualization of the energy footprint history, as shown in Figure 5. The occupant also has the option of setting an alarm threshold, which triggers when the occupant's energy footprint exceeds the threshold.

Figure 6: Energy Footprint Dashboard

5.4.3 Wearables. In addition to the Android and iOS smartphone clients, we develop an application for Android Wear as shown in Figure 5. The Android Wear displays the same energy footprint breakdown that appears in the smartphone client, as well as notifications when alarms are triggered to provide more immediate feedback to the occupant.

5.4.4 Energy Dashboard. Finally, we provide a dashboard at the front of the monitored section of the building that displays various energy consumption statistics, as shown in Figure 6. The statistics shown include energy consumption on a minute-level scale, consumption over the week, consumption breakdown, and an estimated energy bill for the day. Additionally, the dashboard shows the energy footprints of the occupants with the top three highest consumption. The idea is to show to the general public an overview of the general consumption pattern for the building, as well as to reveal the occupants with the most consumption in hopes of motivating occupants to save energy (with user consent).

6 SYSTEM EVALUATION

6.1 Deployment Setup

To evaluate ePrints, we conducted a 2.5 month experiment inside a commercial building to collect personal energy footprints. The testbed area, which consists of 2 full floors inside of a 20-story campus building, is shared by multiple departments. The testbed area is 10,000 square feet, covering diverse types of spaces including offices, cubicles, wet labs, and conference rooms. A total of 53 BACNet endpoints, 15 lights, and 29 plugmeters were monitored.

We invited occupants in the area to participate in our research study. We gathered energy footprint data from 22 participants over the span of 10 weeks. Participants installed either our iOS or Android application, and registered for ePrints.

6.2 Apportionment Policy

The system should be flexible enough to handle different space level and individual level policies. Figure 7 shows the energy footprint of an occupant when different policies are applied. The starting situation is shown in Figure 7a. Halfway through the day, occupant C leaves, leaving space A unoccupied. Figure 7b show the "traditional" equal division policy [8], where the consumption is divided

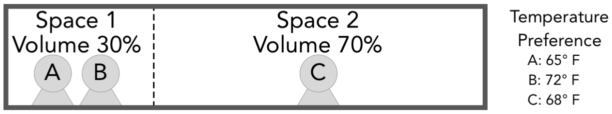

(a) Apportionment policy starting scenario. Space 1 and Space 2 are in the same HVAC zone.

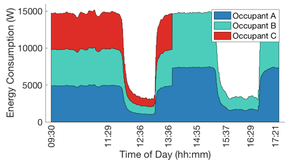

(b) No space policy, equal division individual policy.

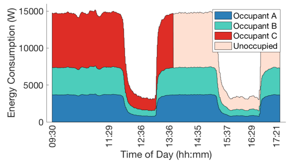

(c) Equal division space policy, equal division individual policy.

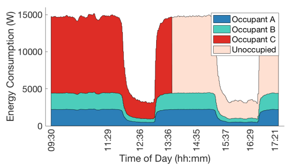

(d) Proportional volume space policy, equal division individual policy.

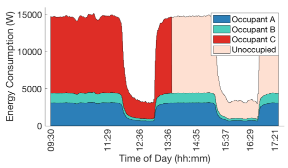

(e) Proportional volume space policy, proportional isolation individual policy.

Figure 7: Apportionment of an energy consuming resource using different space and individual policies.

evenly over the people in the space. When occupant C leaves, his apportioned energy is distributed over occupants A and B.

Figure 7c shows the same situation, but with an additional space policy (equal division). In 7d, the space policy is changed to proportional volume; occupants A and B share an area 30% of the total space; thus, they each claim 15% of the total energy consumption. Finally, 7e changes the individual policy to proportional isolation (defined in [26], with occupant A's isolation consumption at 3000 W and occupant B's isolation consumption at 7000 W).

We also use the criteria defined in [8] to evaluate the aggregate policy defined in equation 1. The first criterion is completeness, the idea that the sum of apportioned energy to individuals should result in the total energy consumed. It can be shown that, if space apportionment policy f and individual apportionment policy g satisfy the completeness criterion, then the comprehensive apportionment policy h also satisfies completeness.

Proof. *Given:*
$\sum_{s\in S} f(s,d) = 1$, $\sum_{p\in P} g(p,s,d) = 1$, $h(p,s,d) = f(s,d)g(p,s,d)$.
Prove: $\sum_{s\in S}\sum_{p\in P} h(p,s,d) = 1$.
$\sum_{s\in S}\sum_{p\in P} h(p,s,d) = \sum_{s\in S}\sum_{p\in P} f(s,d)g(p,s,d)$.
$\sum_{s\in S} f(s,d)\sum_{p\in P} g(p,s,d) = 1$ □

The second criterion, accountability, is the idea that the actions by one occupant should maximally affect their own energy consumption while minimally affecting others. It can be seen from Figure 7 that combining a space-level policy with an individual-level policy separates the groups of occupants. In Figure 7b with the common "equal division" individual policy, the actions of occupant C detrimentally affect occupants A and B. On the other hand, in Figures 7c, 7d, and 7e, the actions of occupant C have no effect on occupants A and B due to the inclusion of the space level policy.

The addition of space-level policy to apportionment maintains completeness and improves accountability for situations in commercial buildings, as in 7a. Thus, this model serves as an improvement to other apportionment models in commercial buildings.

6.3 Zoning

As described in Section 4.1.1, the partition of the total space into zones is critical, as an improper partition of the space may result in apportionment policies that cannot be implemented in ePrints. Figure 8 shows five partitioning strategies: 1) division of the total space using building floor plans, 2) division using HVAC zones, 3) division using organization ownership, 4) division by occupant consensus, and 5) division by automatic space clustering using the method defined in Section 4.1.1. We partition by occupant consensus, by polling occupants about how the spaces should be partitioned, to use as a comparison against the automatic clustering method. The final three methods allow for space-level apportionment policies, such that large energy resources that influence multiple spaces can be divided based on space characteristics.

By measuring the area of overlap between the partitions arising from the automatic space clustering method and from occupant consensus, we see that the two methods yield an 81% similarity value, showing that the clustering method presented in Section 4.1.1 can output similar zones as when space ownership is known.

6.4 Energy Footprint Computation Scalability

In Section 4.2, we introduced two sets of algorithms, one which propagates footprint updates immediately, and one which postpones updates until a footprint query. There are multiple characteristics of the deployment that affect the runtime of these algorithms, including the branching factor (average number of devices affecting a space) and the frequency of each type of event.

To determine the best set of algorithms for this task, we simulated different deployments running ePrints by varying the deployment size and the frequency of occupant location changes, energy resource changes, and footprint queries. The simulations were performed on an Intel 2.4 GHz Core i7 processor. We made a few realistic assumptions based on our testbed: two people and four energy consuming resources per human-centric zone, and based on [6], each person is allocated close to 250 square feet within a commercial building. Finally, we made the baseline assumptions that on average, each occupant is responsible for a footprint query every 2 seconds, changes location once a minute, and energy consuming resources report their energy consumption every 30 seconds.

The point at which a system can not handle the number of queries is denoted by a straight red line in Figure 9. In Figure 9a, we show the effect of the deployment area on the runtime of the footprint computations. As the deployment area increases, the number of occupants and the computation time of every type of event also increases linearly. The propagation algorithms outperform the delayed update algorithms as the deployment area increases. The same can be seen in Figure 9b; starting with the baseline, as we increase the number of footprint queries while maintaining the other parameters, the propagation algorithms perform better than the delayed update algorithms. Since the advantage of the propagation algorithms are lower latency of footprint queries, these results are

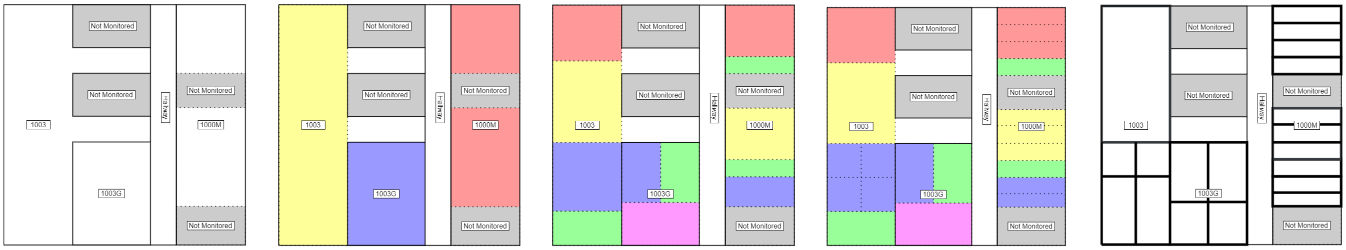

(a) Division by building floorplan schematic. **(b) Division by HVAC zones.** **(c) Division by space ownership groupings.** **(d) Division by occupant consensus.** **(e) Division by automatic space clustering.**

Figure 8: Partition of the testbed using different zoning methods.

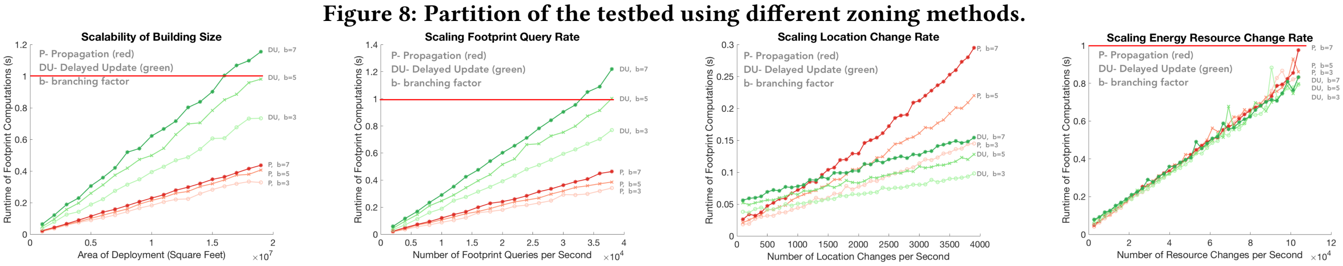

(a) Runtime of footprint computation algorithms with varying deployment area. **(b) Runtime of footprint computation algorithms with varying footprint query rate.** **(c) Runtime of footprint computation algorithms with varying location change rate.** **(d) Runtime of footprint computation algorithms with varying resource change rate.**

Figure 9: Scalability and comparison of "propagation" and "delayed update" algorithms.

intuitive. Also note that the propagation algorithms can support a deployment of 20 million square feet in area and branching factor of 7 at less than 50% processor utilization.

In Figure 9c, we again begin with the baseline assumptions, and increase the number of location changes. As the frequency of location changes increases, the delayed update algorithms outperform the propagation algorithms; however, it should be noted that such a deployment would require a large amount of occupant location changes to take advantage of the delayed update algorithms.

Finally, in Figure 9d, we start with the baseline assumptions and increase the number of energy resource changes. This deployment may include resources that have a high energy reporting rate or more resources than the initial assumptions. As the frequency of resource changes increases, the delayed update algorithms outperform the propagation algorithms by a margin of $10 - 20\%$.

Based on the results of the simulations, a deployment resembling our testbed would benefit from implementing the propagation algorithms. Further, the propagation algorithms are scalable to larger deployments, and can support different environments with varying branching factors, number of occupants, number of energy consuming resources, and number of human-centric zones.

7 ENABLED APPLICATIONS

7.1 Energy Saving Suggestions

One application ePrints enables is energy saving suggestions based off of personal energy footprints. By monitoring a footprint along with major events (such as occupancy changes, HVAC consumption changes, or location changes), one can create a table of events along with each event's average energy saving and the number of occurrences as shown in table 2.

Event Name	Average Energy Saved	Event Occurrence
Move from Lab B to Workspace B	7.1 kW	98
Move from Lab A to Workspace B	5.9 kW	148
Move from Lab A to Workspace A	2.1 kW	170
Occupancy increase in Lab A	2.0 kW	332
Occupancy increase in Workspace A	204 W	38
VAV node A temperature change	296 W	4
VAV node B temperature change	250 W	2
VAV node C temperature change	139 W	7

Table 2: Energy Saving Events

The table of events can be used to inform energy saving suggestions, how effective an occupant's action is in terms of energy savings, and which actions are more likely to be taken by an occupant. For example, a move from Lab A to Workspace B decreases energy consumption by 5.9 kilowatts of energy on average. If the occupant historically has taken this action only a few times, it is unlikely that suggesting this action will be in line with the occupant's intentions. On the other hand, a move from Lab A to Workspace A may decrease energy consumption by fewer watts on average, but occurs more often; thus, this suggestion may be more effective.

Without personal energy footprints, informing energy saving suggestions is difficult if not impossible. Suggestions can be optimized in various ways by analyzing personal energy footprints.

7.2 Building Level Policies

Personal energy footprints are valuable for buildings that allow zone-based HVAC and lighting control, by providing insights into potential actions, and have to potential to save significant amounts of energy when paired with building level policies.

An indirect consequence of the personal energy footprint, along with the space definitions, is the ability to directly derive space level occupancy and energy consumption, and hence the unapportioned energy consumption (energy consumption in unoccupied spaces).

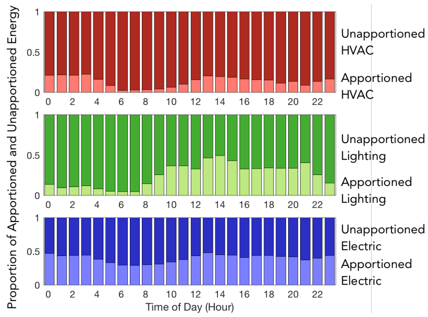

Figure 10: Energy consumption utilization over time, for different energy consuming resource types.

In our deployment, a number of spaces are not only unused during the nighttime, but also unused during the daytime. Figure 10 shows the average proportion of unapportioned and apportioned energy of all spaces in our deployment, over a 10 week span. During the nighttime, there is low energy utilization (high proportion of unapportioned energy) as most spaces are unoccupied. However, it is also notable that during the daytime, a significant portion of energy consumption remains unapportioned; this can be attributed to certain large spaces being almost completely unoccupied throughout the day. This knowledge can help building managers design building level policies and promote energy saving by sensing and optimizing spaces with low energy utilization.

8 CONCLUSIONS

In this paper, we first present design challenges for energy footprinting in commercial buildings. We propose a novel space-centric policy for fair apportionment of energy in shared environments and demonstrate a method for automatically determining space-centric energy zones. We design and implement ePrints – a system for tracking personal energy consumption in real-time. ePrints supports different apportionment policies, with μs-level footprint computation time and graceful scaling with size of building, frequency of energy updates, and rate of occupant location changes. Finally, we presented applications enabled by ePrints, such as mobile and wearable applications to provide users timely feedback on the energy consequences of their actions, as well as energy saving suggestions and building level policies.

ACKNOWLEDGMENTS

The authors would like to thank Columbia Facilities for providing access to the building management system.

REFERENCES

[1] G. Anastasi, R. Bandelloni, M. Conti, F. Delmastro, E. Gregori, and G. Mainetto. 2003. Experimenting an indoor bluetooth-based positioning service. In *Distributed Computing Systems Workshops, 2003. Proceedings. 23rd International Conference on.* IEEE, 480–483.

[2] Bharathan Balaji, Xu Jian, Anthony Nwokafor, Rajesh Gupta, and Yuvraj Agarwal. 2013. Sentinel: occupancy based HVAC actuation using existing WiFi infrastructure within commercial buildings. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems.* ACM, 17:1–17:14.

[3] Bharathan Balaji, Hidetoshi Teraoka, Rajesh Gupta, and Yuvraj Agarwal. 2013. ZonePAC: Zonal Power Estimation and Control via HVAC Metering and Occupant Feedback. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings.* ACM, 1–8.

[4] Nirupama Bulusu, John Heidemann, and Deborah Estrin. 2001. Adaptive beacon placement. In *Distributed Computing Systems, 2001. 21st International Conference On.* IEEE, 489–498.

[5] Sudarshan S Chawathe. 2008. Beacon placement for indoor localization using bluetooth. In *Intelligent Transportation Systems, 2008. ITSC 2008. 11th International IEEE Conference on.* IEEE, 980–985.

[6] Microdata File. 2015. Commercial buildings energy consumption survey (CBECS). (2015).

[7] Emilio Garcia-Fidalgo and Alberto Ortiz. 2015. Vision-based topological mapping and localization methods: A survey. *Robotics and Autonomous Systems* 64 (2015), 1 – 20.

[8] Simon Hay and Andrew Rice. 2009. The Case for Apportionment. In *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings.* ACM, 13–18.

[9] P3 International. 2017. Energy Saving Solutions, Solar Devices and More. http://www.p3international.com/. (2017).

[10] Ruoxi Jia, Ming Jin, Zilong Chen, and Costas J Spanos. 2015. SoundLoc: Accurate room-level indoor localization using acoustic signatures. In *Automation Science and Engineering (CASE), 2015 IEEE International Conference on.* IEEE, 186–193.

[11] Xiaofan Jiang, Stephen Dawson-Haggerty, Prabal Dutta, and David Culler. 2009. Design and Implementation of a High-fidelity AC Metering Network. In *Proceedings of the 2009 International Conference on Information Processing in Sensor Networks.* ACM, 253–264.

[12] Xiaofan Jiang, Chieh-Jan Mike Liang, Feng Zhao, Kaifei Chen, Jeff Hsu, Ben Zhang, and Jie Liu. 2011. Demo: Creating Interactive Virtual Zones in Physical Space with Magnetic-induction. In *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems.* ACM, 431–432.

[13] Yifei Jiang, Xin Pan, Kun Li, Qin Lv, Robert P. Dick, Michael Hannigan, and Li Shang. 2012. ARIEL: Automatic Wi-fi Based Room Fingerprinting for Indoor Localization. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing.* ACM, 441–450.

[14] Andreas Kamilaris, Balaji Kalluri, Sekhar Kondepudi, and Tham Kwok Wai. 2014. A literature survey on measuring energy usage for miscellaneous electric loads in offices and commercial buildings. *Renewable and Sustainable Energy Reviews* 34 (2014), 536 – 550.

[15] Younghun Kim, Thomas Schmid, Zainul M. Charbiwala, and Mani B. Srivastava. 2009. ViridiScope: Design and Implementation of a Fine Grained Power Monitoring System for Homes. In *Proceedings of the 11th International Conference on Ubiquitous Computing.* ACM, 245–254.

[16] Seungwoo Lee, Daye Ahn, Sukjun Lee, Rhan Ha, and Hojung Cha. 2014. Personalized Energy Auditor: Estimating personal electricity usage. In *2014 IEEE International Conference on Pervasive Computing and Communications (PerCom).* IEEE, 44–49.

[17] Joshua Lifton, Mark Feldmeier, Yasuhiro Ono, Cameron Lewis, and Joseph A. Paradiso. 2007. A Platform for Ubiquitous Sensor Deployment in Occupational and Domestic Environments. In *Proceedings of the 6th International Conference on Information Processing in Sensor Networks.* ACM, 119–127.

[18] A. Markham, N. Trigoni, D. W. Macdonald, and S. A. Ellwood. 2012. Underground Localization in 3-D Using Magneto-Inductive Tracking. *IEEE Sensors Journal* 12, 6 (2012), 1809–1816.

[19] N. T. Nguyen, R. Zheng, and Z. Han. 2013. UMLI: An unsupervised mobile locations extraction approach with incomplete data. In *2013 IEEE Wireless Communications and Networking Conference (WCNC).* IEEE, 2119–2124.

[20] Shwetak N. Patel, Thomas Robertson, Julie A. Kientz, Matthew S. Reynolds, and Gregory D. Abowd. 2007. 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.* 271–288.

[21] Ling Pei, Ruizhi Chen, Jingbin Liu, Tomi Tenhunen, Heidi Kuusniemi, and Yuwei Chen. 2010. Inquiry-based bluetooth indoor positioning via rssi probability distributions. In *Advances in Satellite and Space Communications (SPACOMM), 2010 Second International Conference on.* IEEE, 151–156.

[22] Dane Peterson, Jay Steele, and Joe Wilkerson. 2009. WattBot: a residential electricity monitoring and feedback system. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems.* ACM, 2847–2852.

[23] A. Sheinker, B. Ginzburg, N. Salomonski, L. Frumkis, and B. Z. Kaplan. 2013. Localization in 3-D Using Beacons of Low Frequency Magnetic Field. *IEEE Transactions on Instrumentation and Measurement* 62, 12 (2013), 3194–3201.

[24] Juan D Tardós, José Neira, Paul M Newman, and John J Leonard. 2002. Robust mapping and localization in indoor environments using sonar data. *The International Journal of Robotics Research* 21, 4 (2002), 311–330.

[25] Shailja Thakur, Manaswi Saha, Amarjeet Singh, and Yuvraj Agarwal. 2014. WattShare: Detailed Energy Apportionment in Shared Living Spaces Within Commercial Buildings. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings.* ACM, 30–39.

[26] Ekhiotz Jon Vergara, Simin Nadjm-Tehrani, and Mikael Asplund. 2016. Fairness and Incentive Considerations in Energy Apportionment Policies. *ACM Trans. Model. Perform. Eval. Comput. Syst.* 2, 1 (2016), 2:1–2:29.

[27] X. Zhao, Z. Xiao, A. Markham, N. Trigoni, and Y. Ren. 2014. Does BTLE measure up against WiFi? A comparison of indoor location performance. In *European Wireless 2014; 20th European Wireless Conference.* 1–6.