

The Case for Apportionment

Simon Hay
sjeh3@cam.ac.uk

Andrew Rice
acr31@cam.ac.uk

Computer Laboratory
University of Cambridge

Abstract

Apportioning the total energy consumption of a building or organisation to individual users may provide incentives to make reductions. We explore how sensor systems installed in many buildings today can be used to apportion energy consumption between users. We investigate the differences between a number of possible policies to evaluate the case for apportionment based on energy and usage data collected over the course of a year. We also study the additional possibilities offered by more fine-grained data with reference to case studies for specific shared resources, and discuss the potential and challenges for future sensor systems in this area.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous

General Terms

Measurement, Economics

Keywords

Personal Energy Meter

1 Introduction

As part of the Computing for the Future of the Planet research theme [10], we are investigating the concept of a Personal Energy Meter [7]. We envisage a system that collects information about an individual's daily consumption (direct and indirect) and provides breakdowns of the energy costs of our activities to help us target areas for reduction in our environmental footprint [12].

There is evidence to suggest that providing real-time feedback on energy consumption lead to significant reductions [15, 16]. In 2004, buildings accounted for 37% of total energy consumption in the EU [2]; however, economists warn of 'grave inefficiencies' resulting from scenarios where bills are split evenly without regard for individual consumption as

each person minimises their own losses by taking advantage of others [5]. It is this phenomenon that encourages people to order the most expensive items from the menu when out for dinner with a group of friends: if the final sum is to be divided evenly, nobody wants to be subsidising his fellow diners. The same is true of energy consumption in shared buildings: in a house of four where all bills are split, the marginal cost to any individual of turning on an appliance is only a quarter of what it would otherwise be.

Sensors offer the potential to change this balance and apportion energy costs to those who cause them to be incurred: the person standing at the photocopier should be responsible for the energy it consumes during that period, and the cost of the electricity required by a television should be split between all those watching it. There are many challenges which must be overcome in order to achieve an appropriate level of sensor coverage to provide this information.

In this paper we investigate the apportionment of the electricity consumption of our office building. We infer occupancy data from security access logs and show how the choice of policy can have a big impact on personal energy bills. We go on to iteratively refine these initial estimates through the addition of further sensing. We conclude by describing our plans for addressing the challenges in wide-scale sensing for apportionment.

2 Apportionment Policies

We define apportionment as the process of dividing up the total consumption of a building, organisation or other entity and allocating it to individuals. There are numerous possible policies to determine how this should be carried out and different policies will suit different buildings and organisations. Nevertheless, there are certain desirable properties that all apportionment policies should exhibit: 1) Completeness: the sum of the energy apportioned to all individuals should be equal to the total energy to be apportioned and 2) Accountability: actions by an individual should have a maximal effect on their own allocation and a minimal effect on others

The result of apportionment is necessarily specific to a particular individual and so we consider three representative individuals for our building: a member of staff working a standard 9-5 day, a PhD student who arrives later but works the same number of hours and a visiting professor who works part time and has a long commute. Details are shown in Fig-

	Description	Pattern	Hours
1	Member of staff	0900-1700 Mon–Fri	40
2	PhD student	1100-1900 Mon–Fri	40
3	Visiting professor	1100-1700 Tue,Thu	12

Figure 1. Working patterns of example individuals

	Person 1	Person 2	Person 3
Equal	150	150	150
Occupants	132	107	28.9
Occupants+base	168	160	135
Personal load	160	160	143
Personal load+print	168	160	143

Figure 2. Total energy (kWh) allocated by the apportionment policies for a week in November 2007

ure 1. We now develop a variety of apportionment policies and evaluate them with respect to these three individuals for a typical week in November 2007. A summary of the total apportioned energy for each policy is given in Figure 2; these policies are explained in more detail in the following sections.

3 Static Apportionment

The most obvious policy is simply to apportion a static fraction of the building’s power consumption to all those who work there. The number of people allocated desks in our building is around 250 and total energy consumed in 2008 was 2025778 kWh, meaning a user allocated a $\frac{1}{250}$ share would be responsible for 8103 kWh. For comparison, the total energy consumption for one author’s house for the same year was around 2200 kWh.

The electricity meter of our office building, in common with those of many large buildings, logs half-hourly measurements of the total energy consumed. The resulting power apportioned for our example week is shown in Figure 3. The line is a scaled version of the overall power consumption; all users pay more on weekdays, regardless of whether or not they were present. This policy violates our principle of accountability by making no accommodation for working patterns or individual actions—any and all consumption is shared amongst all building users.

4 Dynamic Apportionment

Dynamically varying the proportion of the building’s consumption assigned to each individual allows us to capture the variation in energy due to individual activity. In this first instance we assume that all building users behave similarly and so perform apportionment based on the current occu-

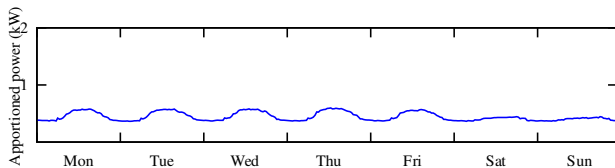


Figure 3. Power apportioned to each individual under the ‘equal’ policy

pancy of the building. A variety of sensor systems could be used to provide this information, including fully fledged location systems [9], existing building access control mechanisms and second order information such as computer activity. Clearly, dedicated sensors provide the best quality data, but we are unlikely to see widespread adoption of these technologies (with their own associated energy consumption) solely to improve energy metering. It is therefore interesting to investigate how we can make use of systems that are in place today before adding more sensing.

4.1 Estimating Occupancy

Although some parts of our building such as the café and lecture theatres are open to all during the normal working day, access cards are required to access most of the office space or to gain entry to the building outside office hours. Holding a card up to a reader unlocks the door from the outside; from the inside, a green button releases it to let people out. The security system keeps logs of all the ‘entry’ and ‘exit’ events and identifies each user on entry with a pseudonym that changes each day. Since multiple people can enter or leave for a single unlock if someone holds the door open, and the identity of those leaving is not determined, we cannot use these to infer who is in the building at any given time. However, the logs can provide us with a reasonable estimate of the overall occupancy. Many buildings have similar systems, but use gates instead of doors and require users to swipe out as well as in; clearly, the records from these systems would be ideal for our purposes.

If we were to assume that one person enters or leaves for each logged event, the running estimate of the occupancy of the building would rapidly drop below zero since, in general, there are approximately 1.25 ‘exit’ events logged for each ‘entry’ event. In order to maintain a stable estimate of the building population we use the following algorithm

1. Count the total number of distinct pseudonyms in a 24 hour period, and assume this is the maximum occupancy for that day (this will under-count people who only entered while someone else held the door open, but it will also over-count because not everyone seen in a day will necessarily have been in the building at once);
2. Calculate the ratio between people entering on ‘entry’ events and people leaving on ‘exit’ events so that the occupancy drops to zero at 5 AM (the logs show this is statistically the quietest time);
3. Scale each day’s estimates so that the peak occupancy is equal to the total number of ‘entry’ events calculated in step 1.

The estimates this produces for our example week are illustrated in Figure 4. Occupancy and power usage are strongly correlated; the occupancy drops off dramatically at weekends, and dips at lunchtime are clearly visible.

4.2 ‘Occupants’ policy

Our first dynamic policy is to split the instantaneous power consumption amongst only those individuals who are in the building at the time. The results of this policy for our example week are shown for several typical working patterns

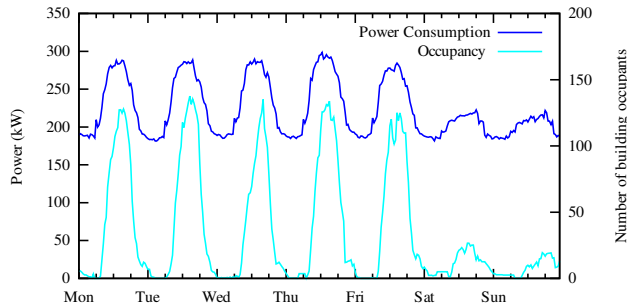


Figure 4. Correlation between estimated occupancy and power usage for the William Gates Building

as the dark lines in Figure 5. There is significant variation dependent on working hours: the example visiting Professor (bottom graph) has a small allocation, but this policy penalises the staff member who now sees large spikes early in the morning when few people are in. This is because the building exhibits cyclic load: many lights and other devices operate on timers or are triggered by movement detectors, so as soon as a few people arrive in the morning the load jumps. In fact this policy strongly discourages any use of the building at unusual times (but this might be the goal). Critically, however, our principle of completeness is violated in that the sum of the energy allocated to all the individual users is not necessarily equal to the total energy consumed by the building: if nobody is in, no energy is apportioned.

4.2.1 Base load

To improve on this policy, we can estimate the base load and divide this amongst all those who work in the building before splitting the remaining power amongst the actual occupants. The results of this calculation are also shown as the pale lines in Figure 5. The base load is estimated as the lowest power consumption seen so far that day. As expected, the peaks during the working day are lower, and the graph no longer drops to zero when a person leaves, instead reflecting his share of the ongoing base load. The sum of the energy apportioned is now equal to the total energy consumed, so from this point of view this policy represents an improvement. Intuitively, the policy is better, too, since now all those who have reserved office space in the building are held responsible for some share of its ongoing costs.

The graphs still display several peculiarities. In particular, two people working the same number of hours are allocated substantially different amounts of energy because fewer people are in at 9 AM than at 11 AM but a large proportion of the shared energy consumers (lighting etc.) have already been switched on.

The policy also runs into problems with accountability: if the base load is shared amongst all users of the building while the additional energy consumed is divided between the occupants at the time, it is in an individual's best interests to maximise the base load (of which he is only allocated a small fraction). One way to exploit this policy is by leaving his computer and lights on overnight—this results in a *lower* energy cost to him than switching them off, since they are then included in the base load and split between many more

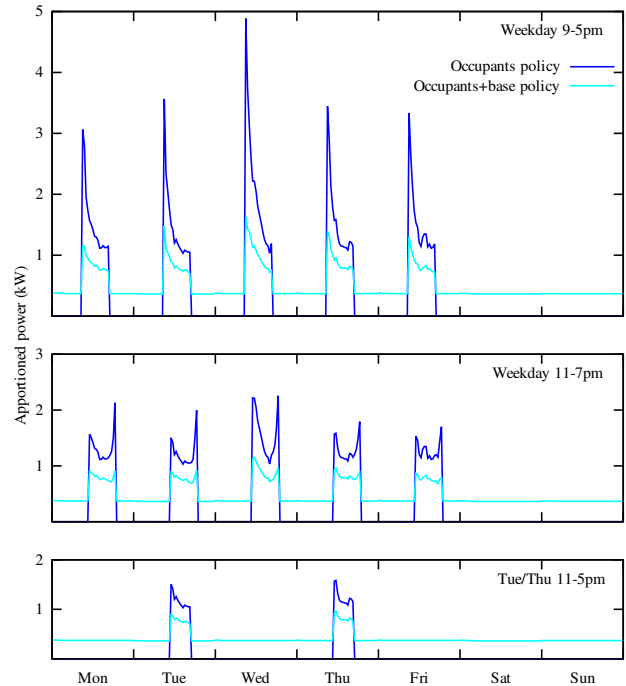


Figure 5. Power apportioned under the 'occupants' policies to example individuals

people.

4.3 'Personal load' policy

Instead of estimating the base load and assuming the remainder is personal, we can approach the problem from the opposite direction by estimating the personal load and assuming the remainder should be divided evenly. The 'personal load' policy allocates a certain amount of power to each occupant of the building and then divides the remainder evenly amongst all users. A survey of one of our offices with a simple power meter revealed that the devices we all typically switch on when we arrive (lights, monitors) consume between 100 and 200 W, depending on office size and computer configuration. Supporting this observation, our dataset reveals that 150 W is a sensible average figure to allocate to each occupant—any more results in the total energy allocated to occupants dipping beneath our earlier estimate of the base load. Figure 6 shows the results of this policy for the same three sample individuals as before.

The output of this policy is reminiscent of the previous one, as we might hope and expect, but the incentives now work in the correct direction: the motivation for an individual is to do his best to reduce his own energy consumption. For these incentives to work the effect of any changes made must be visible in the results, and this entails a more detailed measurement of power consumption than we have considered so far. Instead of simply dividing up the total energy bill for the building, it will be necessary to identify which specific devices an individual uses and how much power they all require.

Clearly, the different policies make a significant difference to the end results, but it is interesting to note that

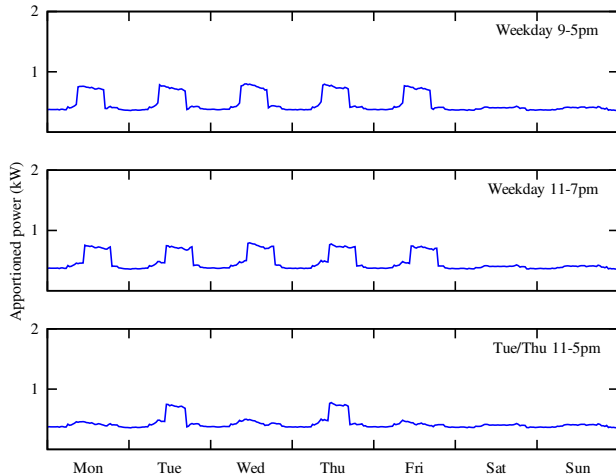


Figure 6. Power apportioned under the ‘personal load’ policy to example individuals

the ‘occupants+base’ and ‘personal load’ policies produce broadly comparable numbers. This supports our intuition that both are reasonable strategies and the only difference between them is a result of an inaccurate estimate of the base load: if we had omniscient sensor data and could tell exactly which devices were consuming power the two policies would become the same.

5 Refining Personal Load

The accountability of the ‘personal load’ policy can be incrementally improved by adding particular information about the components of each user’s personal load. In certain cases this is easy: we define the notion of *owned resources*, which belong in some sense to an individual who is in general held responsible for their energy consumption. Someone with a private office might be considered to own everything in it; certainly everyone will ‘own’ the computers and monitors on their desks.

Many resources do not fit into the pattern described above; they are not owned, but *communal*, shared by a group or team. Printers, photocopiers, projectors, coffee machines and showers fall into this category. The obvious mechanisms for handling these resources are either simply to divide their total energy cost amongst all those entitled to use them or to attempt to ascertain who is using them at any given time and allocate the energy used accordingly.

The energy used must now be measured at a much finer resolution; we require additional sensors that can measure the usage of a corridor, room or specific device. Previously knowing the cumulative energy consumed was sufficient and the required update frequency was dictated only by the desired reporting period; to apportion energy costs based on usage, we must measure the energy consumed in each individual interaction. As a middle ground short of continuous online measurement, we built custom hardware that integrates the readings of an off-the-shelf clamp meter and logs the timestamped results card at 10 Hz. This high frequency allows us to identify the energy costs of specific events, such as printing a page or making a cup of coffee. The intention

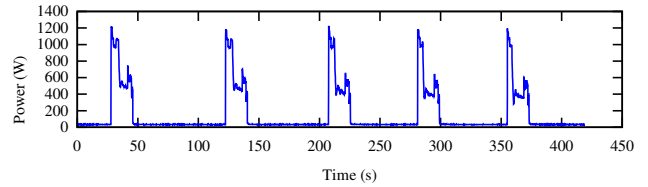


Figure 7. Power drawn printing five single pages

here is to profile a device in detail and then use some other indication of usage to infer energy consumption.

5.1 Printing

Printing provides a good case study of the value of apportioning use of shared resources since we have second order information on usage in the form of print server logs.

We performed an analysis of the printer logs for our building. The logs cover 28 printers for a period of 47 days, during which time 198 users printed a total of 82349 pages. During this period there were 313 users with accounts who were entitled to use the printers.

The data shows a large deviation between different users’ printing habits. The heaviest user printed 3452 pages, while the lightest printed just 1, and the top 15 users accounted for over half the total printing between them.

As an illustration, we measured the energy consumption of one printer using the apparatus described above. An example trace is shown in Figure 7. The printer draws 32 W when idle (in power save mode), and consumes, on average, an additional 11200 J to print a single page. Printing multiple pages at once costs less per page than printing a single page on its own as the warm up costs are amortised; we found that the average energy cost per page for the whole workload over several days was 8720 J. Assuming these figures to be typical of all printers, the average energy cost per day of having the printers switched on was 21.5 kWh, with an additional 4.24 kWh consumed by printing.

We can now combine our energy measurement with the logs from the print server to improve on the ‘personal load’ policy. The energy consumed by a particular print job was originally spread over all occupants and so for each job we remove a share of the energy consumed from all occupants before reassigning the total to the individual who printed the material. The results of this policy for our chosen staff member, who printed a large set of lecture course material in the week in question, are shown in Figure 8 and represent an increase of 8 kWh for the week. Neither of our other example users printed anything; their results show a reduction in allocated energy of around 0.5 kWh for the week.

5.2 Anonymous Shared Resources

Our research group shares a coffee machine. In this case, as for most shared resources, there is no second order information available on its usage. We are left with several possible options that require varying investments of time and infrastructure. We could ignore their usage entirely and allow them to be counted as part of the base load; we could ask users to keep track of their usage manually; we could deploy some form of identity prompt, such as the PINs often required on photocopiers for accounting purposes, or we

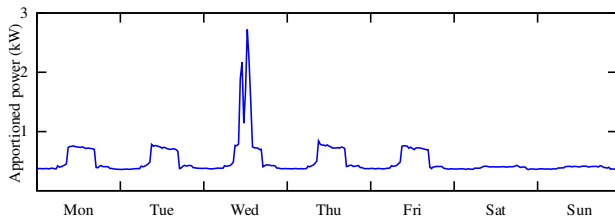


Figure 8. Apportionment with printing costs

could install a separate sensor system. The most appropriate will depend on the significance of the resource in question.

Our measurement apparatus showed that the electrical energy required to make a single cup of coffee is approximately 62 KJ, or 0.02 kWh, and two cups coffee per day accounts for about 3% of an individual’s 150W personal load. This is totally insignificant relative to the power draw of the whole building and it would not be worth the energy cost of a sensor system to apportion its use. Ignoring it altogether would probably be justifiable—but lessons learnt from the coffee machine can be applied to all sorts of other equipment.

To evaluate the feasibility of the manual method, during the course of one week we asked members of our research group to make a mark against their name on a tally sheet every time they had a coffee. We also included an ‘Anonymous’ row to allow those who preferred not to have their usage recorded to participate in the study.

25 separate people logged their consumption, ranging from only 1 cup in the whole week to 17. For comparison, there were 53 registered members of the research group or visiting students during the week in question. Out of 212 cups of coffee logged, 58, or 27%, were anonymous. However, the machine’s own audit trail shows that in fact 333 cups of coffee were produced over the period in question; only 64% of cups were logged. This suggests that a number of people chose not to record their usage on grounds other than privacy concerns, even for research purposes when no attempt at charging was being made—most probably on account of the extra time and effort involved. We conclude that any attempt to apportion the use of these resources as part of a future personal energy meter must therefore be entirely unobtrusive and automatic, requiring no additional intervention on the user’s part; schemes such as RFID readers that require swiping a access card, or logon systems, will probably irritate users and not be adopted unless they are made compulsory (i.e. integrated with the appliance itself). This is unlikely to be practical in the majority of real world situations.

Location systems promise to provide all the input required for accurate apportionment, revealing exactly who is in a building at any given time and (generally) who is using a particular device (although depending on the resolution of the system ambiguities may remain where a number of people are gathered around). Figure 9 shows the trace of a user walking from his office to the coffee machine recorded using the Bat system [1] — but note that since no sensors are installed in the kitchen, it is ambiguous exactly which device

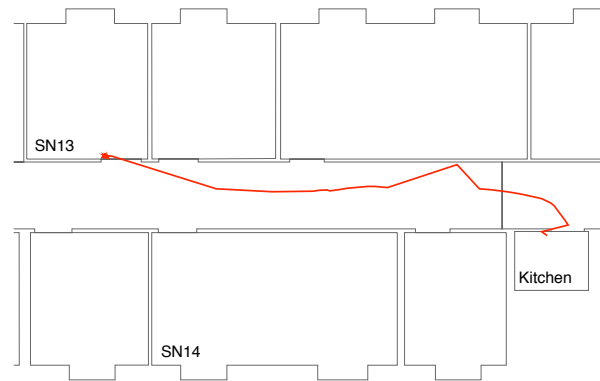


Figure 9. Location trace of walking to the coffee machine

he was using. Matching locations against usage logs may reveal the user in most cases, but even with perfect instrumentation, it will be difficult to distinguish which of several people next to a machine is using it, or whether he is doing so on behalf of someone else.

The majority of these systems have not spread outside research labs, primarily due to the cost of deploying and maintaining building wide location technologies. Large amounts of custom hardware must be deployed, surveyed and calibrated; users must all remember to wear an additional device, and in many existing environments the infrastructure requirements are simply impractical. It seems unlikely that institutions will choose to deploy such a system solely for the purpose of energy apportionment.

The ideal solution seems to be *context awareness* in devices, so they themselves were aware of the identity of the user and so could personalise their response accordingly. A coffee machine might recognise an RFID tag in a user’s mug and so produce his preferred drink automatically: this would serve as a ‘carrot’ providing a valuable incentive rather than a ‘stick’ forcing users to comply [14].

Any viable system for our purposes must have a very low cost, both in monetary, infrastructure and energy terms. One technology we have investigated that has the potential to meet these requirements is tracking based on Bluetooth [8], but there are several other possibilities, and we believe this is an exciting and relevant area of research.

6 Scaling up

While the owners of certain classes of equipment such as computers may be well known, maintaining an up-to-date inventory, alongside either detailed power profiles for devices or continuous monitoring sensors, is a challenge. One possible solution is to encourage users to maintain their own data on devices they own or supply additional sensor data: providing additional or more accurate information voluntarily could result in a reduction in apportioned energy. This technique has seen success with water companies encouraging home owners to fit meters. Device profiles and even office layouts could be shared, allowing everyone to benefit from and build on the work of others. Establishing such commu-

nities around the data may have another benefit: social networks provide an ideal forum for users to share consumption patterns and reduction strategies.

7 Related Work

Energy apportionment cuts across a broad range of research areas, from energy monitoring through location and identity sensing systems to human-computer interaction and social questions. Krumm et al. have used sensors that detect electrical noise on power lines and machine learning techniques in an attempt to recognise use of certain electrical equipment [11]. In the field of detecting a user's presence, Harle has investigated the potential for using location systems to optimise energy consumption dynamically [6], while Garg and Bansal show how to improve on the estimates of simple occupancy sensors by adapting to changing activity levels [4]. Dodier et al. explore the use of belief networks with occupancy sensors [3]. There is a significant body of work on simulating occupancy profiles using Markov chains where live data is unavailable [17, 18]. There has also been research on presenting this data to end users: in particular, Mankoff et al. explore how social networks can motivate users to reduce their ecological footprints [13].

8 Conclusions and Future Work

Through the simulation of several policies for example individuals, we have shown that apportionment is important and the correct choice merits careful consideration. Different policies have significant effects on the total energy allocated to individuals. We believe that personal load provides the best opportunity to personalise results and improve accuracy incrementally and offers valuable incentives for users to reduce their consumption.

Although reasonable estimates can be made from sensor data commonly available today, more precise analysis requires investment of time in power profiling and inventory management as well as research into novel, low-infrastructure identity and location sensing systems. This represents a building block towards a true Personal Energy Meter.

9 Acknowledgments

We would like to thank Andy Hopper for his insight and support, Martyn Johnson and Chris Hadley for their help in providing access control and printer logs and Rob Harle, Alastair Beresford and Dan Ryder-Cook for their advice and suggestions.

10 References

- [1] M. Addlesee, R. Curwen, S. Hodges, J. F. Newman, P. Steggle, A. Ward, and A. Hopper. Implementing a sentient computing system. *IEEE Computer*, 34(8):50–56, Sep 2001.
- [2] I. E. Agency. Key world energy statistics. 2006.
- [3] R. Dodier, G. Henze, D. Tiller, and X. Guo. Building occupancy detection through sensor belief networks. *Energy and Buildings*, 38(9):1033–1043, Sep 2006.
- [4] V. Garg and N. K. Bansal. Smart occupancy sensors to reduce energy consumption. *Energy and Buildings*, 32(1):81–87, Jun 2000.
- [5] U. Gneezy, E. Haruvy, and H. Yafe. The inefficiency of splitting the bill. *The Economic Journal*, 114(495):265–280, 2004.
- [6] R. Harle and A. Hopper. The potential for location-aware power management. *UbiComp '08: Proceedings of the 10th International Conference on Ubiquitous Computing*, Sep 2008.
- [7] S. Hay. A global personal energy meter. *Adjunct Proceedings of the 7th International Conference on Pervasive Computing*, 2009.
- [8] S. Hay and R. Harle. Bluetooth tracking without discoverability. *LoCA 2009: Proceedings of the 4th International Symposium on Location and Context Awareness*, 2009.
- [9] J. Hightower and G. Borriello. Location systems for ubiquitous computing. *Computer*, Jan 2001.
- [10] A. Hopper and A. Rice. Computing for the future of the planet. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 366(1881):3685–3697, Oct 2008.
- [11] J. Krumm, G. D. Abowd, A. Seneviratne, and T. Strang. At the flick of a switch: Detecting and classifying unique electrical events on the residential power line. *Lecture Notes in Computer Science: UbiComp 2007: Ubiquitous Computing*, 4717, Jan 2007.
- [12] P. Kurp. Green computing. *Communications of the ACM*, 51(10), Oct 2008.
- [13] J. Mankoff, D. Matthews, S. Fussell, and M. Johnson. Leveraging social networks to motivate individuals to reduce their ecological footprints. *System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference on*, pages 87 – 87, Jan 2007.
- [14] K. Mansley, A. R. Beresford, and D. Scott. The carrot approach: Encouraging use of location systems. *UbiComp '04: Proceedings of the Sixth International Conference on Ubiquitous Computing*.
- [15] J. E. Petersen, V. Shunturov, K. Janda, G. Platt, and K. Weinberger. Does providing dormitory residents with feedback on energy and water use lead to reduced consumption? *Proceedings, Greening The Campus VI*.
- [16] J. E. Petersen, V. Shunturov, K. Janda, G. Platt, and K. Weinberger. Dormitory residents reduce electricity consumption when exposed to real-time visual feedback and incentives. *International Journal of Sustainability in Higher Education*, pages 16–33, 2007.
- [17] I. Richardson, M. Thomson, and D. Infield. A high-resolution domestic building occupancy model for energy demand simulations. *Energy and Buildings*, 40(8):1560–1566, Jan 2008.
- [18] D. Wang, C. Federspiel, and F. Rubinstein. Modeling occupancy in single person offices. *Energy and Buildings*, 37(2):121–126, Feb 2005.