Technical Report

UCAM-CL-TR-805
ISSN 1476-2986

Number 805

UNIVERSITY OF CAMBRIDGE

Computer Laboratory

A model personal energy meter

Simon Hay

September 2011

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500

http://www.cl.cam.ac.uk/

This technical report is based on a dissertation submitted
August 2011 by the author for the degree of Doctor of
Philosophy to the University of Cambridge, Girton College.

Technical reports published by the University of Cambridge
Computer Laboratory are freely available via the Internet:

http://www.cl.cam.ac.uk/techreports/

ISSN 1476-2986

A model personal energy meter

Simon Hay

Abstract

Every day each of us consumes a significant amount of energy, both directly through transport, heating and use of appliances, and indirectly from our needs for the production of food, manufacture of goods and provision of services. This dissertation investigates a personal energy meter which can record and apportion an individual's energy usage in order to supply baseline information and incentives for reducing our environmental impact.

If the energy costs of large shared resources are split evenly without regard for individual consumption each person minimises his own losses by taking advantage of others. Context awareness offers the potential to change this balance and apportion energy costs to those who cause them to be incurred. This dissertation explores how sensor systems installed in many buildings today can be used to apportion energy consumption between users, including an evaluation of a range of strategies in a case study and elaboration of the overriding principles that are generally applicable. It also shows how second-order estimators combined with location data can provide a proxy for fine-grained sensing.

A key ingredient for apportionment mechanisms is data on energy usage. This may come from metering devices or buildings directly, or from profiling devices and using secondary indicators to infer their power state. A mechanism for profiling devices to determine the energy costs of specific activities, particularly applicable to shared programmable devices is presented which can make this process simpler and more accurate. By combining crowd-sourced building-inventory information and a simple building energy model it is possible to estimate an individual's energy use disaggregated by device class with very little direct sensing.

Contextual information provides crucial cues for apportioning the use and energy costs of resources, and one of the most valuable sources from which to infer context is location. A key ingredient for a personal energy meter is a low cost, low infrastructure location system that can be deployed on a truly global scale. This dissertation presents a description and evaluation of the new concept of inquiry-free Bluetooth tracking that has the potential to offer indoor location information with significantly less infrastructure and calibration than other systems.

Finally, a suitable architecture for a personal energy meter on a global scale is demonstrated using a mobile phone application to aggregate energy feeds based on the case studies and technologies developed.

Acknowledgments

A great many people have willingly given their time and effort to help me with my Ph.D., to whom I offer my sincere thanks and my regrets that there is insufficient space to provide a comprehensive list.

In particular, I thank Andy Hopper for his vision, insight and support, and Alastair Beresford, Robert Harle and Andy Rice for their advice, guidance, suggestions and assistance. I could not have managed without their help, and I remain extremely grateful.

I am indebted to Brian Jones, whose patience and knowledge has been invaluable. I also wish to thank George Coulouris, Joe Newman and Ian Wassell for their direction and feedback which helped shape my research.

Thanks to Dan Ryder-Cook, who created the physics model described in Section 4.2 and David Piggott, who implemented the Android application described in Section 6.2. Thanks to Oliver Woodman for developing the original version of OpenRoomMap on which a number of my measurements depend; he, Ripduman Sohan and my other officemates and members of the Digital Technology Group have provided a friendly and encouraging atmosphere in which it has been a pleasure to work.

Richard Bird developed my interest in Computer Science at Lincoln College, Oxford; without his suggestion and tuition I would not have been able to contemplate a Ph.D. Sandy and Elisabeth Fraser supported me and broadened my education. I am also grateful to my examiners, Adrian Friday and Peter Robinson, whose feedback has markedly improved this dissertation.

Finally, I wish to thank Claire and my parents, family and friends for their love and support.

Publications

Some of the contributions presented in this work have appeared in the following peer-reviewed journal, conference and workshop publications:

- Simon Hay and Robert Harle. Bluetooth tracking without discoverability. In *Proceedings of the 4th International Symposium on Location and Context Awareness (LoCA 2009)*, Tokyo, Japan. DOI: 10.1007/978-3-642-01721-6_8
- Simon Hay and Andrew Rice. The case for apportionment. In *Proceedings of the 1st ACM Workshop On Embedded Sensing Systems For Energy-Efficiency In Buildings (BuildSys 2009, in conjunction with SenSys 2009)*, Berkeley, CA, USA. DOI: 10.1145/1810279.1810283
- Andrew Rice and Simon Hay. Decomposing power measurements for mobile devices. In *Proceedings of the 8th IEEE International Conference on Pervasive Computing and Communications (PerCom 2010)*, Mannheim, Germany. *Winner of the Mark Weiser Best Paper Award.* DOI: 10.1109/PERCOM.2010.5466991
- Andrew Rice and Simon Hay. Measuring mobile phone energy consumption for 802.11 wireless networking. *Pervasive and Mobile Computing*, Volume 6, Issue 6, December 2010, Pages 593–606, ISSN 1574-1192. DOI: 10.1016/j.pmcj.2010.07.005
- Andrew Rice, Simon Hay and Dan Ryder-Cook. A limited-data model of building energy consumption. In *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys 2010, in conjunction with SenSys 2010)*, Zurich, Switzerland. DOI: 10.1145/1878431.1878447

In addition, some content has been taken from the following reports, conference abstracts and position papers:

- Simon Hay, Andrew Rice and Andy Hopper. A global personal energy meter. *Ubiquitous Computing at a Crossroads Workshop: Art, Science, Politics and Design (UbiComp Grand Challenge)*, London, UK.
- Simon Hay, Joseph Newman and Andrew Rice. Sentient computing meets social networking. *W3C Workshop on the Future of Social Networking*, Barcelona, Spain.
- Simon Hay. A global personal energy meter. In *Adjunct Proceedings of the 7th International Conference on Pervasive Computing (Pervasive 2009)*, Nara, Japan.
- Simon Hay, Andrew Rice and Andy Hopper. Personal energy metering. *CompSust '10: The 2nd International Conference on Computational Sustainability*, Cambridge, MA, USA.

Finally, the following conference and workshop publications arose from other work not presented in this dissertation:

- Simon Hay, Joseph Newman and Robert Harle. Optical tracking using commodity hardware. In *Proceedings of the 7th IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR 2008)*, Cambridge, UK.
  DOI: 10.1109/ISMAR.2008.4637345
- Stamatina Th. Rassia, Simon Hay, Alastair Beresford and Nick Baker. Movement dynamics in office environments. In *Proceedings of the 3rd CIB International Conference on Smart and Sustainable Built Environments (SASBE 2009)*, Delft, Netherlands..
- Simon Hay, Stamatina Th. Rassia and Alastair Beresford. Estimating personal energy expenditure with location data. In *Proceedings of the 1st IEEE Workshop on Pervasive Healthcare (PerHealth 2010, in conjunction with PerCom 2010)*, Mannheim, Germany. DOI: 10.1109/PERCOMW.2010.5470650
- Salman Taherian, Marcelo Pias, Robert Harle, George Coulouris, Simon Hay, Jonathan Cameron, Joan Lasenby, Gregor Kuntze, Ian Bezodis, Gareth Irwin and David Kerwin. Profiling Sprints using On-Body Sensors. In *Proceedings of the 6th IEEE International Workshop on Sensor Networks and Systems for Pervasive Computing (PerSeNS 2010)*, Mannheim, Germany. DOI: 10.1109/PERCOMW.2010.5470629

Press

The following media stories relate to work presented here:

- The Tech Lab: Andy Hopper. *BBC News*, 23 April 2009. `http://news.bbc.co.uk/1/hi/technology/8014248.stm`
- John Walko. Researchers ready personal energy monitoring devices. *EE Times*, 17 June 2009. `http://www.eetimes.com/electronics-news/4195584/Researchers-ready-personal-energy-monitoring-devices`
- Where next for the web? *BBC News*, 9 March 2010. `http://news.bbc.co.uk/1/hi/technology/8555987.stm`
- Discovery—Superpower: The Future of the Internet. *BBC World Service*, 14 March 2010. `http://www.bbc.co.uk/programmes/p006hrrg`
- Wendy M. Grossman. Using technology to reduce the carbon footprint. *The Inquirer*, 7 April 2010. `http://www.theinquirer.net/inquirer/feature/1599632/using-technology-reduce-carbon-footprint`

Contents

List of figures

List of tables

Chapter 1

Introduction

Information is differences that make a difference. (Edward Tufte)

Contents

Overview

This chapter introduces the research theme of *Computing for the Future of the Planet* and sets out the vision and the need for a personal energy meter. It explains the principal benefits and discusses the challenges, both technical and social, involved in realising such a project on a global scale. It outlines a strategy for tackling these problems and states the research questions addressed in the remainder of this dissertation.

1.1 A personal energy meter

Every day each of us consumes a significant amount of energy, both directly through transportation, heating and use of appliances, and indirectly from our needs for the production of food, manufacture of goods and provision of services. Controlling the use of natural resources will be one of the world's greatest challenges in the years to come; whether for reasons of climate change, scarcity, economy or limiting dependence on foreign powers, reducing our energy requirements has undeniable benefits [119].

An answer to this problem might lie in the use of pervasive computing technologies to build a personal energy meter: a device which can record and allocate an individual's energy usage in order to provide baseline information and incentives for reducing the environmental impact of our lives [97]. The idea of a personal energy meter has garnered some media attention (see page 9). It might help suggest areas for improvement and, if an individual makes a lifestyle change, analyse its impact and significance [128]. This would depend on a global sensor network and poses a number of challenges: new sensor systems are required both to account for the energy used and to determine the identity and activities of users. This dissertation develops the principles and technologies necessary to build a model system.

Huge imbalances currently exist between the environmental footprint of individuals in different countries: taking the mean over a year, the average person in North America consumes around 12,000 W, compared with around 6,000 W in Europe, 1,000 W in India and only 300 W in Bangladesh.[1] The global average is around 2,000 W; this can be considered a sustainable target for us each to strive to achieve. The *Société à 2,000 Watts* project calls for a reduction in energy needs to an average of 2,000 W per person.[2] This target considers not only personal or household energy use, but the total for our entire society, divided by the population. The project has gained significant government backing, with a pilot region in Basel and the cities of Zurich and Geneva announcing their intentions to become 2,000 W societies by 2050.

This and similar proposals have attracted significant media attention. For example, in an article for *The Sunday Times* a journalist described his attempt at 'the low-watt diet'.[3] Despite taking all manner of seemingly extreme measures he ultimately encountered the perennial problem—it is at present very difficult for an individual to assess whether or not he is achieving an energy target:

> I went to a dinner party in a part of London where everyone seems to drive a 4×4. I sat next to a woman who listened politely as I described the steps I'd taken towards a 2,000 W life. She wondered if I'd hit the target. Honestly, I had no idea…

To judge the success of a weight-loss programme one requires weighing scales, but all most people have to judge their energy consumption today is infrequent, coarse-grained and incomparable billing information.

The idea of 'footprints' has caught the public imagination; they are readily understood and much-discussed. A number of websites have sprung up recently claiming to calculate a user's personal carbon footprint. For example, AMEE[4] calls itself 'the world's energy meter' while Encraft[5] offers a carbon footprint calculator that takes into consideration household energy bills, private and public transport mileage and flights. These services are obviously very limited in scope and require manual data input, but demonstrate the keen interest in monitoring energy consumption.

In reality most estimates of carbon emissions or ecological areas are simply energy consumption figures scaled by a predetermined factor for the type of energy used and divided equally amongst a large population. Energy seems the best metric because it is easily

[1] World Resources Institute. EarthTrends. `http://earthtrends.wri.org/`

[2] `http://www.societe2000watts.com/`

[3] John-Paul Flintoff. Energy: How low can you go? *The Sunday Times*, 23 November 2008. `http://www.timesonline.co.uk/tol/news/environment/article5188314.ece`

[4] `http://www.amee.cc`

[5] `http://www.encraft.co.uk`

quantifiable and does not attract the controversy that sometimes surrounds estimates of carbon emissions or ecological footprint, but the underlying sensing and calculation techniques are transferable. More careful apportionment can be applied to data, regardless of how it is collected, and offers more meaningful results.

A personal energy meter that provides live information on consumption apportioned to individuals represents a significant step forwards from the common situation of a static, approximate and time consuming audit of a building or organisation. This idea fits into the *Computing for the Future of the Planet* framework for identifying ways in which computing can have a positive effect on our lives and the world [98]; one such way is to *sense* the world around us to inform us about the energy consumption and other effects of our activities on the natural environment. It is dependent on developments in a number of computing technologies—in particular, sensors and sensor networks to provide data both on usage and on interactions and a common world model to allow information to be collected wherever the user might be. Contextual information, such as might be provided by location systems, will be important to determine how the energy costs of shared resource should be apportioned.

1.2 Potential benefits and drawbacks

Personal energy data will enable us to identify areas for optimising our consumption of resources. It will make metering 'smarter' by providing itemised breakdowns for individuals rather than buildings. Projections of consumption will allow us to see the total cost or benefit of a decision to replace an appliance, install insulation or move house. The personal energy meter will also make offsetting schemes more realistic and help us identify alternatives to our current activities. For example, the trace of commuting to work might be analysed to highlight any suitable public transport available or to inform policy for providing future facilities. Many energy reducing measures also bring monetary savings, but there is insufficient information available to let consumers realise their best courses of action. A personal energy meter could help its users make more informed decisions.

Aggregate measures may raise awareness, but apportioning consumption to individuals will be crucial to helping us identify where changes could most effectively be made and if those changes have had any meaningful impact. A user seeing the breakdown of his share of an office building's consumption may realise that significant savings could be made by switching his computer off overnight; if every user in the Computer Laboratory made this change the department's electricity bill could be reduced by around £6,100 each year (Section 4.2.4). However, the impact of one individual making such a change would be lost in the noise of a plot of the building's consumption and so users might be discouraged from taking action. Personal data might even help bring about improvements that are outside any individual's control: when users have made all the changes in their power they may be frustrated to find that they are still allocated significant energy costs for items such as heating resulting from poor insulation and so collectively lobby for higher-level (or even political) changes.

1.3 Architecture

The design and implementation of a personal energy meter embodies many challenges. Input data must come from a wide range of meters, sensors and systems distributed globally; calculations can take place either locally or remotely and there are several forms of output from a personal counter to shared statistics. Effective communication with a planetary 'world model' must be maintained in order to provide up-to-date estimates of energy consumption. In this sense, the personal energy meter fits within *sentient computing*, which Hopper defines as "using sensor and resource status data to maintain a model of the world which is shared between users and applications" [96].

There are three main parts to personal energy metering:

1. data collection
2. processing
3. presentation

Figure 1.1 illustrates some of the data flows envisaged. Feedback could be presented using the same data in several separate forms, such as via a website, on a mobile device or even using shared public displays [205]. A model architecture for disseminating data is proposed in Chapter 6, while Chapters 4 and 5 describe mechanisms for obtaining the consumption and context information required.

1.4 Research questions

This dissertation tackles the issue of how sentient computing can apportion the energy costs of shared resources to individuals and addresses four main research questions:

To what extent can technology be used to apportion personal energy costs? What would be needed that does not exist already?

One of the most interesting challenges of a personal energy meter is in apportioning the energy costs of large shared resources such as office buildings and public transport to individuals. Context awareness offers the potential to apportion energy costs to those who cause them to be incurred, which may provide incentives to make reductions. Careful thought should be given to the correct policies for apportionment and to what inputs these policies require. To what extent is live sensing necessary?

Can energy consumption be inferred without continuous metering? Are there classes of energy consumer that must be treated differently?

A key ingredient for apportionment mechanisms is data on energy usage. Although more and more smart meters are being deployed, it is unrealistic to expect the usage of every individual device to be monitored in real time in the near future. Instead, can models, inventories and profiles be used to estimate consumption?

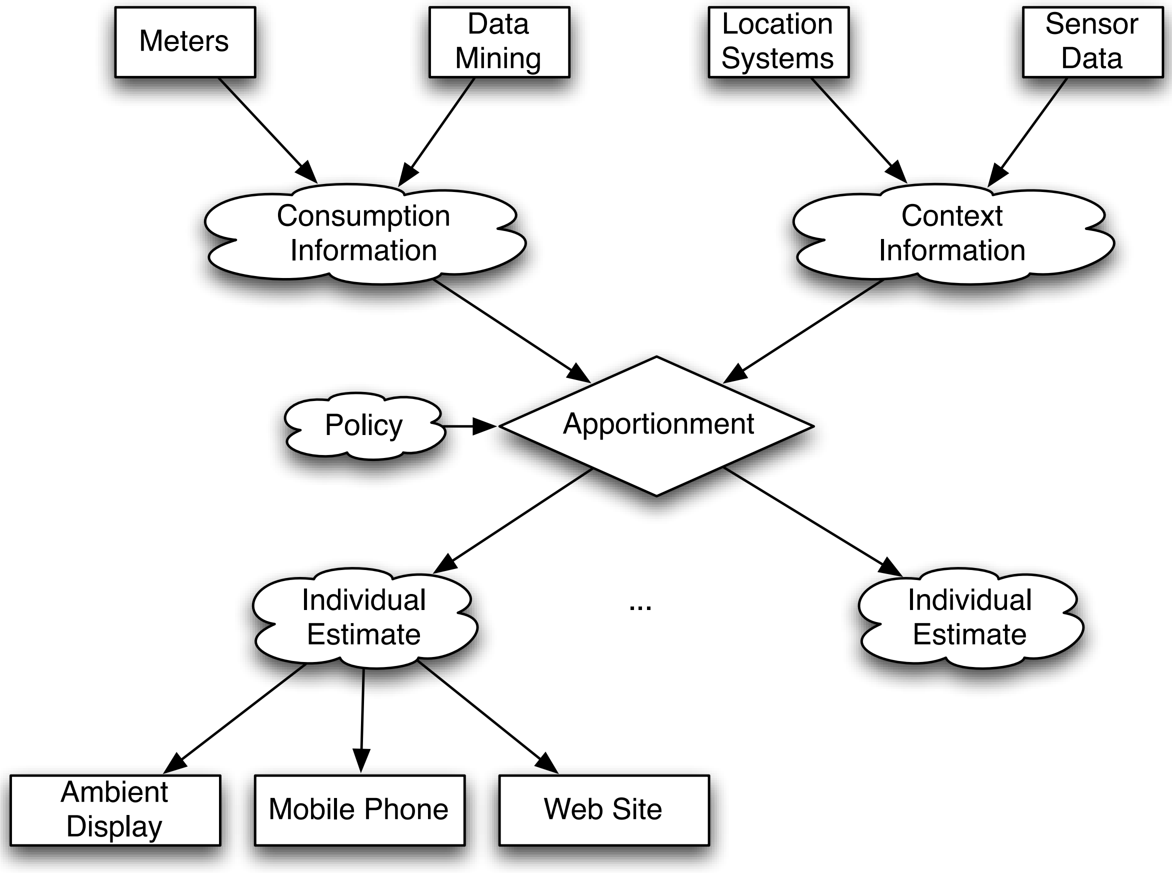

Figure 1.1: Example personal energy meter architecture

Can context be monitored with minimal additional infrastructure? How can the benefits of location systems to provide context for apportionment be obtained without the related costs?

Contextual information provides crucial cues for apportioning the use and energy costs of resources (Section 5.1), and one of the most valuable sources from which to infer context is location. Indoor location systems have been the subject of much research over the past two decades, but while many systems can deliver impressive results very few are suitable for widespread deployment outside research environments due to the extensive bespoke infrastructure that must be installed and surveyed. This is costly in terms of both money and time, and impractical in most buildings. A key ingredient for a personal energy meter is a low cost, low infrastructure location system that can be deployed on a truly global scale. Can the necessary information be obtained by repurposing hardware that is already deployed?

What should be the software architecture of a personal energy meter? How can it be made to scale to global proportions, and how could users be encouraged to adopt it?

To piece together the complete picture of his energy usage, each user will require information from many separate sensor systems which each meter individual parts of his overall consumption. How should this information be aggregated?

Although the future may be a world of sensors, it is important to consider how to build a system that scales gracefully from today's sparse and unreliable sensing through to a vision of total knowledge. Most users are unlikely to be prepared to install a large number of complicated sensors solely in order to use a personal energy meter. Throughout, this dissertation therefore considers the problem of deployability of sensors, and how much information can be obtained with as little sensing as possible through the use of humans as sensors and repurposing other pieces of infrastructure.

1.5 Dissertation outline

The structure of the remainder of this dissertation is as follows:

Chapter 2 reviews related work in relevant areas of research. Existing techniques for influencing behaviour, measuring energy consumption and inferring context are described.

Chapter 3 presents the results of two studies to motivate the remainder of this dissertation. It describes the importance of apportioning energy costs to individuals, evaluates a range of strategies and derives the key ingredients required: *metering* and *context*.

Chapter 4 investigates how sentient computing systems can meter or calculate energy consumption, and the specific sensing requirements of shared resources. It presents a technique for using crowd-sourced inventories and device profiles to estimate building energy consumption and explores to what extent it is possible to use existing ground truths to infer or estimate the consumption of devices not covered by sensor systems. Finally, it describes a novel framework for decomposing power measurements of programmable devices necessary to apportion energy costs of specific actions to individuals.

Chapter 5 identifies the key characteristics required of a location system to provide context for a personal energy meter. It presents and evaluates a novel mechanism for inquiry-free Bluetooth tracking that has the potential to provide the low-cost, pervasive tracking necessary.

Chapter 6 examines a model architecture for scaling a personal energy meter to planetary proportions. It proposes a federated system of energy feeds and demonstrates it with a mobile phone application that aggregates and visualises the data from systems described in previous chapters.

Chapter 7 concludes by revisiting the research questions posed in Section 1.4, outlining possible avenues for future research and summarising the main contributions of this dissertation.

1.6 Limitations of scope

The construction of a truly global personal energy meter applicable to anyone, anywhere, would be an engineering challenge well beyond the scope of a dissertation such as this.

Instead, it addresses mainly the first two challenges set out in Section 1.3 of data collection and processing; since the HCI aspects and exact form in which feedback is presented makes a significant difference to its effectiveness (see Section 2.1) only example interfaces are developed. It focuses on the principles and concepts that are transferable from proofs of concept and case studies centred on the author and colleagues at the Computer Laboratory of the University of Cambridge. It identifies and draws together existing work which might help account for significant sources of energy consumption such as transport, and tackles only those areas which are lacking, with a particular focus on the apportionment of the energy costs of shared resources.

Chapter 2

Related work

Contents

Overview

This chapter provides a detailed survey of related work. The idea of a personal energy meter cuts across a broad range of research areas, from energy monitoring through location and identity sensing systems to human-computer interaction and social questions. This chapter first motivates and places in context the personal energy meter through a review of studies showing the effectiveness of feedback on reducing energy consumption (Section 2.1), then discusses existing technologies designed to promote behavioural change (Section 2.2). It surveys systems for measuring or inferring energy consumption that might provide useful input to a personal energy meter (Sections 2.3 and 2.4), highlights the problems with requirements for extensive additional infrastructure and reviews potential solutions in the form of user-deployed sensing or crowd-sourced data (Section 2.5). Location will provide important context information to help apportion consumption, and Sections 2.7 and 2.8 provide a detailed survey of existing systems, focussing on the essential properties for energy metering. Finally, Section 2.9 discusses potential methods for aggregating data from many disparate sensor systems which will be necessary for any heterogeneous personal energy meter. This chapter therefore identifies where further research is needed and guides the remainder of this dissertation.

2.1 The importance of feedback

Behavioural and environmental psychology studies have demonstrated many times the impact that feedback can have on encouraging people to reduce their energy consumption; this provides strong motivation for the creation of a personal energy meter which provides continuous fine-grained feedback across all the aspects of a person's life rather than on specific places or types of consumption. The first known study of 'eco-feedback' was in 1976, when Kohlenberg et al. found that even a light bulb which illuminated when households were within 90% of their peak energy levels changed energy usage behaviour, reducing time spent above a predetermined peak power level by up to 50% [121]. Real-time web-based feedback has been shown to produce an overall 32% reduction in electricity use by dormitory residents, high resolution feedback proving more effective than low resolution [167]. Feedback is not only useful in its own right, as a self-teaching tool, but it also improves the effectiveness of other information in achieving better understanding and control of energy use. This section highlights some of the more significant reviews to obtain aggregate measures of the effects of feedback.

In a review of 38 feedback studies carried out over a period of 25 years, Darby found typical energy savings of 10–15% and showed that improved feedback may reduce consumption by up to 20%; she concluded that "clear feedback is a necessary element in learning how to control fuel use more effectively over a long period of time" [37].

Fisher surveyed 26 projects from 1987 onward on the effects of feedback of electricity consumption and on consumers' reactions, attitudes and wishes concerning such feedback [53]. She found that typical energy savings were between 5% and 12% and that feedback is most successful when it "is based on actual consumption, given frequently and over a long time, provides an appliance-specific breakdown, is presented in a clear and appealing way and uses computerised and interactive tools." Although only three of the studies reviewed used computerised feedback, these were the ones that resulted in the greatest change in consumption. She hypothesised that successful feedback has to draw a close link between specific actions and their effects; this is one of the core aims of the personal energy meter. Winett et al. also showed the importance of specificity, with more specific signs resulting in a 60% reduction in days lights were left on compared with more general feedback [204].

Abrahamse et al. compared feedback to a number of other intervention strategies, such as goal-setting and information, through a review of 38 studies from social and environmental psychology [1]. They conclude that "feedback has proven its merits, particularly when given frequently" and that it can also increase the effectiveness of antecedent interventions. They argue that an important first step in any intervention aimed at reducing energy is a 'thorough problem diagnosis' to identify the behaviours that significantly contribute to environmental problems; this is one of the key aims of a personal energy meter.

Froehlich et al. published a comprehensive survey of examples of what they call 'eco-feedback technology', including 89 papers from environmental psychology and a further 44 from ubiquitous computing and human-computer interaction literature [63]. They pointed out that although these two fields are closely related they have tended to remain wholly separate and argued that Computer Science researchers have not yet focussed on evaluating the potential strengths of their designs with respect to their ability to change behaviour.

Finally, Fitzpatrick and Smith reviewed the methodology of previous studies of feedback technologies and set out design concerns and questions that they hope will influence future work, focussing on how feedback should be presented [54]. They suggested that current work is too utility-centric and a more holistic view is necessary, imagining that "multiple resources could be monitored, not just electricity and gas, but also water, garbage, chemicals, food, and so on; feedback on light use could be juxtaposed against occupancy and activity monitoring." The personal energy meter attempts to take just such a view.

2.2 Persuasive technologies

The personal energy meter is intended to draw people's attention to specific information in an attempt to change what they do or think. Fogg has labeled this phenomenon 'persuasive technology' and suggested it can be used to change people's behaviour in domains such as preventative healthcare and fitness [58, 57]. Interactivity gives computing a strong advantage over more traditional persuasive media, and ubiquitous computing in particular means that interactive computing systems embedded in everyday objects and environments can intervene at precisely the right place and time to have maximum effect.

Unfortunately, there are very few projects taking an overall view of *personal* consumption as a personal energy meter should. However, Froehlich et al. bemoaned the dearth of feedback technologies in domains other than home electricity consumption and proposed ideas for automatically sensing use of electricity and water, personal transportation, product purchases and garbage disposal and their potential impact for reducing consumption [62]. This overview paper surveyed existing systems and refers to their own ideas, all of which are discussed in more detail in the remainder of this chapter and in Chapter 3. Although they have contributed excellent work on sensor systems and feedback technology, there is not yet a unified system for combining the results of these disparate systems in the manner they suggest; the personal energy meter seeks to fill this gap.

Meanwhile, Mun et al. introduced the idea of a *Personal Environmental Impact Report* [147]. This is a participatory sensing application that uses location data sampled from mobile phones to calculate personalised estimates of environmental impact and exposure, including carbon impact, sensitive site impact, smog exposure and even proximity to fast food restaurants. It highlights the reciprocity between impact and exposure: that which is good for the person is often also good for the planet, and vice versa.

2.2.1 Sensor data communities

Social networks provide an ideal forum for users to share consumption patterns and reduction strategies. There is significant overlap between the fields of sentient computing and social networking; further details are provided in a separate paper [88]. For example, the research literature contains numerous examples which demonstrate the potential of location-based services in wider social networking contexts. One application that is often touted is a 'colleague radar' akin to Harry Potter's 'Marauder's Map' which shows at a glance the locations of coworkers within a building. Examples include the ACTIVEMAP tool [143] and the 'intelligent Coffee Corner' [146]. These have been represented in a variety of forms: the visual similarity between an Active Badge application displaying a simple list of people and locations and a modern RSS feed of friend status is notable

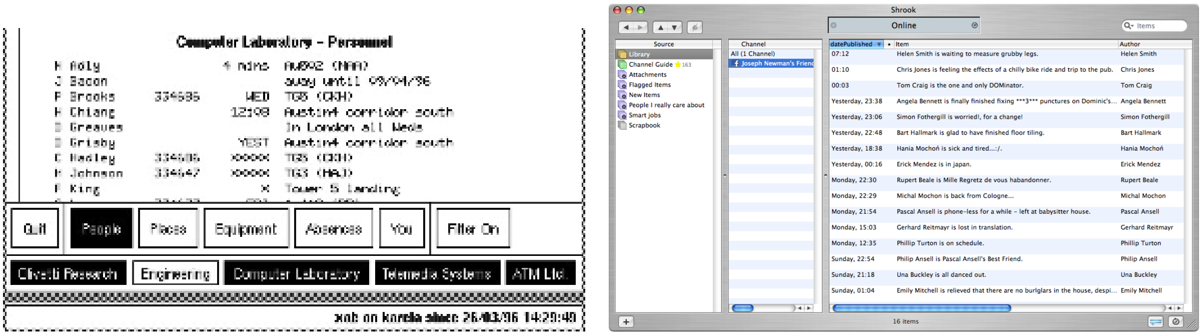

(a) Active Badge (b) RSS feed

Figure 2.1: Updates on friends in 1990 and 2010

(Figure 2.1), and 3D representations of the data resemble metaverses such as Second Life. Location systems provide a way to bridge the physical/digital divide [98].

There are several examples of online communities forming around sensor data to encourage and assist each other. In the context of sport and fitness, many people find it hard to motivate themselves to train as hard or as often as they might like. The Nike+ system[1] automatically captures details of the user's runs from accelerometers in his shoe and can share them on a website, allowing him to define goals, compare his progress to that of others with similar targets, set challenges for friends and exchange training programmes. Dailymile[2] has a similar premise, aiming to make it easier for users to share their workouts, exchange advice and find training partners. This has a number of parallels to energy saving measures, and the same strategies can also be applied (for a fuller discussion, see Section 2.2.2). Dopplr[3] lets its users share estimates of their carbon emissions through travel, and DIY KYOTO has a community site[4] based on its electricity monitoring device, allowing users to share graphs of their power consumption, record their energy reductions and offer money-saving tips to others. It also shows the total amount of energy and money saved by all its users combined, highlighting the global significance of seemingly small scale actions.

Communities have formed around power graphs, sharing advice and experiences. The CenceMe application provides a good example of a mobile application that allows the sharing of context information inferred through sensors via social networks [144]. This includes not only location but also activity: it can detect 'dancing' through the accelerometers, for instance, or 'conversation' through the microphone.

Other related products and systems are surveyed in the remainder of this section. The power of the social networking phenomenon might be brought to bear in a similar way to drive adoption of the personal energy meter and provide impetus and support for changing lifestyles.

[1] http://www.nikeplus.com
[2] http://www.dailymile.com
[3] http://www.dopplr.com
[4] http://community.diykyoto.com/

2.2.2 Persuasive technologies for physical activity

There are significant parallels between encouraging people to exercise more and encouraging them to use less energy: both are things that most people know they should do but find it hard to motivate themselves, and indeed similar sensors can provide useful input data for both. For example, walking to work instead of driving has clear benefits for both the person and the planet; a continuous location trace for an individual might be used both to estimate his energy footprint due to transport and to provide feedback on his personal health based on his physical activity levels. Providing information on personal health might encourage uptake of sensors that also provide data valuable for personal energy metering, and many of the lessons on the efficacy of feedback that have been learned from projects on physical activity are also applicable to energy consumption.

These studies teach that results should be understandable at a glance, show progress towards goals and allow easy comparisons with others or against historical data.

Rogers et al. investigated a range of ambient installations (see Section 2.2.3.1) to encourage building occupants to take the stairs rather than the lift [181]. This approach is rooted in behavioural psychology and the belief that a 'nudge' at the point of decision-making can influence the choice. In *Follow-the-Lights* a display of twinkly white lights was embedded in the carpet, triggered by foot pressure and highlighting a path to the stairwell. *The Clouds* consisted of a set of grey and orange spheres suspended from the ceiling in the atrium; their relative heights changed in relation to the number of people who used the stairs versus the lift that day. Finally, *The History* depicted an aggregate visualisation on a large public display showing historical trends. The authors conducted a number of surveys to gauge opinions on these pieces, and also gathered quantitative data on usage from pressure sensors near the stairs and lifts. During an eight-week study they found that the ratio of stair to lift use increased significantly (from 1.49 to 2.13), highlighting the potential of these feedback techniques.

Chick clique is a mobile phone application that targets a specific demographic, aiming to help motivate teenage girls to exercise by exploiting their desire to confirm to social norms and stay connected with their peers [194]. Each participant can select a group of up to three friends to engage in a friendly competition where the group's walking statistics are tracked. The final design consisted of a pedometer embedded in a belt whose readings were entered manually into an application on the mobile phone, keeping track of the number of steps that were taken each day. Automated text messages should be sent at opportune times indicating the group performance, including individual fitness level achieved, but since the prototype was deployed on a PDA with no cellular connectivity, these also had to be sent manually. The system was tested on two groups of friends and compared against the pedometer alone; while one group did take more steps when given the group information, the other actually walked further without the system.

Fish'n'Steps also attempts to increase daily step count, this time by linking it to the growth and activity of an animated fish [133]. 19 volunteers took part in a 14 week 'Wizard of Oz' study, interacting with a system which they believed to be autonomous but which was in fact partially operated by an unseen human being. Again, the pedometers could not be read automatically, but to simplify the procedure and avoid requiring participants to enter readings, with the attendant potential for tampering with the results, they were asked to place their pedometer on a platform at a public kiosk and take a picture of its pedometer screen, including the unique ID. The picture was sent to a member of the research team who entered the appropriate data into a database. Each participant was set custom goals relative to their baseline, and success in reaching a participants daily goal affected the facial expression of his or her fish. They all apparently found the procedure inconvenient and awkward but participated for the sake of their colleagues' research; the authors admit that interest in the game subsided after a couple of weeks, but claim that it still generated sustainable change in behaviour.

Houston is another application in a similar vein, consisting of a pedometer and a mobile phone application into which users can enter their daily step counts by hand [31]. A three-week study was conducted using three versions of the software: a baseline which merely collected the data, a personal version which also enabled goal setting and tracking progress and a sharing version which allowed participants to send their counts to friends and see friends' progress towards goals. Average step counts increased, with participants using the sharing version significantly more likely to meet their goals.

Building on the lessons from Houston, *UbiFit Garden* is more sophisticated set of persuasive technologies to encourage physical activity [33]. 12 participants in a three-week study wore a *Mobile Sensing Platform* containing a number of sensors including a 3D accelerometer and barometer [28]. The sensing platform runs classifiers trained to distinguish walking, running, cycling, using an elliptical trainer and using a stair machine; the list of activities and their current likelihoods are communicated via Bluetooth to an application running on the user's mobile phone. The application builds a journal of activities which can be added to manually; the user can then review the log and progress towards a weekly goal. A 'glanceable' or ambient display is also included in the background of the phone's home screen, using the metaphor of a garden, with flowers and butterflies representing goal attainments and different types of activity. Feedback was particularly positive towards the idea of the glanceable display; the authors were sufficiently encouraged to conduct a second three-month study with 28 participants focussing on the details of the feedback mechanism [32]. This trial took place over the Christmas period, which is traditionally the hardest time of year for people to keep up exercise regimes; the researchers found that the participants with the glanceable display tended to keep up their level of physical activity while those of the participants without the display dropped significantly. Although the specific representation they chose is clearly not to everyone's taste, they demonstrated that immediate ambient feedback has an important part to play in encouraging behavioural change.

With the exception of UbiFit, none of these projects use sophisticated sensor systems to derive their output; in fact, many rely on manual input (though it is still a valuable lesson to learn that users are prepared to enter data by hand provided the subsequent feedback is sufficiently useful). Their utility is therefore not in ubiquitous computing technologies but in understanding both that feedback can have a persuasive effect and that the way in which it is presented is all-important.

2.2.3 Persuasive technologies for energy consumption

Much of the Computer Science work published on energy consumption has addressed the HCI issues surrounding how best to present feedback to engage users. Although this aspect is outside the scope of this dissertation, which focuses on calculating results rather than their presentation, it is very relevant to the wider concept of a personal energy meter.

Dillahunt et al. attempted to "leverage the power of the polar bear as a symbol of climate change by creating a virtual pet" [44]. Participants were shown a Flash-based virtual polar bear on an ice floe that would grow as they committed to environmentally responsible actions such as 'dry only full loads of clothes' and 'use a low wattage night light' and decrease as they chose not to commit to actions. Users demonstrated significantly greater environmental concern and greater care after reading about and interacting with the polar bear. Those who used the system were found to commit to more actions, fulfil more commitments and donate more to environmental causes than the control group, but closer analysis shows that only one of these differences was statistically significant; the authors pointed out that they could not know how long the differences would be sustained.

Computer games have a natural persuasive power [58]. *PowerHouse* was a prototype computer game that aimed to promote an energy-aware lifestyle among teenagers, employing the format of a reality TV show to inform its users about energy-efficient action [12]. The player managed a simulate domestic environment with the objective of keeping the residents content within a limited energy budget. Unfortunately, no results were presented on its effectiveness as an intervention.

EcoIsland presented a virtual island on a household display, with each family member represented there by an avatar [187]. The family set a target CO_2 emissions level and the system tracked their approximate current emissions using self-reported data. If the emissions exceeded the target level, the water around the island began to rise; carrying out suggested emission-reducing activities caused a drop in water level. It also supported a kind of emissions trading scheme, where reductions could be swapped with neighbours for credit that could be used to decorate the island. In common with many of the persuasive systems described here that relied on self-reporting, a survey showed that after a week of use 17 out of 20 participants felt more conscious of environmental ecology, but there was no statistically significant correlation with actual energy usage.

Mankoff et al. highlighted the potential of social networks for motivating change and outlined a proposal for sharing energy consumption information on sites such as MySpace and Facebook [140]. It is certainly true that social networks have an overwhelming reach and influence and might help create peer pressure to reduce environmental impact. A subsequent paper described StepGreen.org, a website with accompanying plugins for MySpace and Facebook which realised some of these ambitions [139]. It encouraged its users to commit publicly to energy-saving actions which were displayed on their social network profiles both as reminders to themselves and promises to others, and also to compare their progress with that of their friends. Although a small internal user study was carried out which influenced a redesign of the site, little information was provided on its acceptance or success in a wider sense. Nevertheless, social networking could indeed have an important part to play in personal energy metering; this idea is discussed further in Section 2.2.1.

2.2.3.1 Ambient displays

The *Power-Aware Cord* was an augmented electrical power cord which represented the amount of energy passing through it with glowing patterns produced by electroluminescent wires moulded into the transparent cord [75]. The authors argued that electricity, being invisible and intangible, must be made more perceptible to increase awareness of consumer energy consumption and lead people to question their behaviours, and that light is a more natural and intuitive way to symbolise this relationship than a digital readout.

This is an example of ambient display, in the same vein as the famous 'Live Wire', or 'Dangling String' installation by Natalie Jeremijenko at Xerox PARC during Mark Weiser's time there. This was a piece of string attached to a stepper motor and controlled by a LAN connection; network activity caused the string to twitch, yielding a peripherally noticeable indication of traffic. Ambient displays complement the idea of augmented reality in which additional information is added to environmental elements [9]: instead of augmenting the user with a head-mounted display or similar technology physical objects are themselves enhanced.

The *Ténéré* was a project in a similar style, again augmenting a electrical extension cord to provide feedback 'in meaningful and emotional ways' [114]. L'arbre du Ténéré was a solitary acacia in the Sahara that was once considered the most isolated tree on Earth until it was knocked down and replaced with a metal tree-like sculpture. The authors used this as a metaphor for the environmental consequences of human activity. Their device had an OLED display which showed the tree morph into a sculpture if too much energy was consumed.

Holmes presented a public art installation at the National Center for Supercomputing Applications entitled *7000 oaks and counting* [95]. This operated at a building level but also used the tree metaphor, composed of a sequence of animated clips using a series of tree images that corresponded to the carbon loads in the building. It used custom software to gather electricity, condensate, and chilled water usage figures every minute from the building monitoring system and convert the aggregate data to reveal the buildings real time carbon footprint. She called this process 'eco-visualization'. The estimated footprint was in turn converted to the number of trees that would be required to offset the carbon emitted and an animation is displayed which shows greater numbers of trees when the load is very high and a few larger, more detailed trees when it is low. Users can fill out a web form claiming to have offset some of this load and their name is incorporated into the animation and the offset applied to the building's total.

As in Section 2.2.2, few of the papers surveyed here present novel mechanisms for measuring energy consumption; instead, they generally rely on off-the-shelf meters or self reporting. Again, some of the virtual-reality-based feedback methods seem too far-fetched to become popular, but the same results on the importance of feedback and how it is presented that became clear from studies of physical activity are apparent in the literature on energy consumption. This lends further weight to the belief that the two fields are closely related, and lessons learned in one can be applied to the other. Although ambient displays clearly cannot replace more detailed feedback they may complement it well and further raise awareness of consumption; a personal energy meter should therefore support multiple separate feedback mechanisms driven from the same data sources.

2.3 Metering electricity consumption

2.3.1 Direct electricity metering

2.3.1.1 Building level

Building-level metering is becoming increasingly prevalent as utility companies deploy smart meters with machine-readable interfaces. There are also myriad projects, both academic and commercial, aimed at gathering, processing and storing this data which are surveyed in this section. Although building-level consumption does not in itself provide any information on how energy was used, it is a good starting point for a personal energy meter and various techniques can be applied to disaggregate it by both user and function.

Although over time many buildings will transition to smart meters that can report usage automatically, for now manually-entered readings remain vital to avoid excluding a large proportion of the population. ReadYourMeter[5] is a free website that allows users to record and graph their utility meter readings and so aims to help them understand their energy consumption, compare their data with others and see how much energy organisations use. Sites like this will provide a useful source of high-level data for the personal energy meter.

Weiss et al. created a lightweight gateway to capture usage data from meters and expose it through a RESTful API over the web and to users' mobile phones [203]. Their interface showed overall consumption and also helped identify the consumption of individual devices by prompting the user to turn them on and off and measuring the difference. Although the application itself is not especially noteworthy, the two ideas of exchanging sensor data using web technologies and presenting disaggregated consumption on a mobile phone are promising and explored further in Chapter 6.

Patel et al. presented an end-user-deployable, whole house, contactless power consumption sensing system [161]. It consists of two devices: a sensor on the circuit-breaker panel which detects the draw from the magnetic field generated from the 60 Hz current flow, and a plug-in module which cycled through a series of known loads for automatic calibration as well as sensing the line voltage needed to calculate true power. This design has significant advantages: its contactless nature means an end user can simply stick it to the panel using double sided tape, while LEDs illuminate to indicate signal strength and help with positioning; there are no safety concerns. Data is streamed to a PC via Bluetooth at approximately 1 kHz. Experiments in three different homes showed the average error after four weeks was around 4%. The authors then conducted a study to determine ease of deployment; eight participants (chosen from 73 respondents to an online survey) were asked to install the devices in their own homes, and all completed the task successfully in an average of 20 minutes. In the online survey, 86% claimed they would be comfortable installing it on their own. These are very positive results, and this represents the state of the art in user-deployable high-resolution whole-building metering. The same team go on to attempt to detect and classify electrical events; see Section 2.3.2 for a fuller discussion.

Agarwal et al. gathered data on the electricity consumption of four representative buildings on the campus of the University of California at San Diego, configuring meters to report data several times a second to a central acquisition server and presenting it to users in the form of an online Energy Dashboard [7]. They augmented building-level metering with 15 separate circuit-level instruments in one building to provide a more detailed breakdown of the load, discovering that much of it is due to IT infrastructure; other work proposed mechanisms to reduce the base load of computing equipment [5, 6]. Kleissl and Agarwal subsequently conducted further analysis of the sub-metered building and determined that energy savings of around 80% for lighting, 60% for computing, 50% for server rooms and 20% for mechanical loads were possible [120].

As well as research systems there are an ever-increasing number of commercially-available tools for monitoring energy consumption. Any attempt at a comprehensive list would

[5] `http://readyourmeter.org/`

quickly become outdated, but a few examples are highlighted here to demonstrate the range of offerings.

Google PowerMeter is an application being developed by Google's philanthropic arm, Google.org, to help consumers track their home electricity usage.[6] Google has partnered with both utility companies, who provide usage data directly on behalf of their customers who have opted in to the service, and manufacturers of energy monitoring devices such as those described previously. The information is stored on Google's servers and a web interface provides visualisations of home energy consumption and allows users to set targets, track progress and discuss their findings with others in the online community; it supports both normative and historical self-comparison. Data can also be both uploaded and retrieved via a public API.

Microsoft Hohm is an equivalent competing offering from Microsoft, though it is at the time of writing only available in the United States via energy feeds from utility companies rather than standalone meters.[7]

TED (The Energy Detective) is a building-level energy meter which is installed in the breaker panel.[8] It will upload its recorded data to Google PowerMeter as well as hosting a built-in web server which allows remote analysis. Its software also supports load profiling of up to five individual appliances. Current Cost have a similar range of devices that also store and report their data to Google; remote display units are also available.[9] AlertMe offer both meter readers and plug monitors that communicate using ZigBee and report consumption to both display devices and Google.[10]

Wattson and Holmes, a hardware and software combination from DIY KYOTO, aim to bring contemporary design to energy metering and offer an ambient display of consumption through mood lighting. They store 28 days of usage information and allow data and energy-saving hints and tips to be shared in an online community.

2.3.1.2 Circuit level

Taherian et al. described an energy monitoring platform called *Cambridge Sensor Kit Energy* designed to measure usage with off-the-shelf clamp meters at aggregation points such as meters and fuse boards and provide feedback both at the local deployment site and remotely over the Internet [191]. Data processing and aggregation takes place on an embedded device which provides a web server to show graphs and details of consumption as well as an RSS feed and RESTful API to support alternative feedback mechanisms such as a digital photo frame and three-coloured LED globe.

2.3.1.3 Appliance level

There has been a flood of commercially-produced appliance-level monitors, which are now available at very low cost. The 'Kill A Watt' is just one example, displaying volts, amps,

[6] `http://www.google.com/powermeter/`
[7] `http://www.microsoft-hohm.com/`
[8] `http://www.theenergydetective.com/`
[9] `http://www.currentcost.com/`
[10] `http://www.alertme.com/`

watts, Hz, and VA for a single electrical socket.[11] Many of these are not machine readable, showing their measurements only on a built-in display, but their popularity indicates a serious interest amongst the general public in better understanding their electricity consumption.

Energy Optimizers Limited produce the Plogg[12] range of electricity power meters, which are supplied either as stand-alone end user devices or as embedded hardware for incorporation into OEM products. These combined intelligent plugs and data loggers sample the voltage and current signals and store the values for subsequent wireless retrieval over Bluetooth or ZigBee. The Bluetooth versions have proved useful but are let down by software that makes obtaining results in a machine-readable format difficult and unreliable. However, as part of the *Energie Visible* project, Guinard et al. extended the basic functionality of the Ploggs to include continuous reporting and built a gateway that discovered the devices and exposed them to the web through a RESTful API [73]. They deployed the system on a single floor of their office building to monitor the energy consumption of various devices; a large display enables people passing by to experiment with the energy consumption of the devices. They argued that HTTP and web technologies are a good match for sensor nodes, allowing 'mashups' to be created of their data using familiar tools and techniques as part of the 'Web of Things' [72]. This idea is discussed further in Section 6.2.

The *Plug* sensor network, first mentioned in 2006 [160] was a set of 35 sensor-, radio-, and computation-enabled power strips distributed throughout the MIT Media Lab [132]. A single Plug device functioned as a normal power strip but also had range of sensors (sound, light, electrical current and voltage, vibration, motion, and temperature) for gathering data about how it was being used and its nearby environment. This is an interesting approach to the problem of distributing sensors in a building, since it involves replacing existing hardware with smarter alternatives rather than deploying additional equipment; many of the traditional problems faced by wireless sensor networks such as battery life and being as unobtrusive as possible cease to exist. This is closer to the vision of disappearing hardware originally set out by Mark Weiser as 'ubiquitous computing' [199]; the computing is integrated into the infrastructure rather than being a separate artefact.

The power consumption of the monitoring infrastructure itself is also important, especially if it is intended to meter every device in a large building. Bai and Hung designed a board to provide remote power control and current measurement for electric sockets via a ZigBee wireless connection, along with an embedded home server to gather and store the readings [11]. This approach results in lower power consumption from the measuring infrastructure itself, with the sockets using 318 mW and the embedded server 1-3 W (compared to several hundred watts for a conventional PC).

The *Bit-Watt* system used a network of so-called 'smart-taps', which are sockets that sample their voltage and current at 6 kHz and report the data, again via ZigBee, to a home server [131]. The server attempted to identify each appliance by matching the wave shape pattern of the current values against a database of known signatures, although no results were offered on its accuracy or reliability. A 3D visualisation was then presented to the user.

Jiang and a team from the University of California, Berkeley built a wireless sensor net-

[11] `http://www.p3international.com/`
[12] `http://www.plogginternational.com/`

work called *ACme* for monitoring AC energy usage in a large building environment [104]. Each of the nodes, which are based on their previously-discussed TinyOS-based *Epic* architecture, plugs into a single socket and exposes its energy consumption to arbitrary endpoints via a IPv6 interface (they chose to use 802.15.14 as their physical layer rather an 802.11 because of the difficulty of adding new nodes to an existing wireless network).

The nodes sample at 14 kHz, which is sufficient to identify patterns and potentially attempt appliance recognition, but typically report energy readings once per minute via UDP to a simple Python daemon running on a server. Each node draws an additional 1 W of power from the mains socket it is plugged into; this power consumption is low, but not negligible if considered on the scale of an entire building which may have thousands of individual sockets and therefore result in additional load measured in kilowatts. A preliminary deployment lasted 4 months with 49 nodes split across a university department and a private apartment. Its scalability and ease of deployment makes this the best architecture for widespread, scalable device-level power metering presented to date, and it will probably form the basis of developments in the future.

The ACme design is open source, and in a demonstration of its general applicability Kazandjieva et al. modified the design slightly, adding an expansion port to support new sensors and storage and enabling a sealed case, then built and deployed 85 of the meters throughout the Stanford Computer Systems Lab [112, 113]. They are integrated into their *PowerNet* infrastructure, which also includes 55 commercial wired meters; deployment of the latter was subsequently abandoned due to issues of size and proprietary protocols. The aggregate measurements of all the power meters amount to only 2.5% of the building's total consumption, but by compiling a device inventory from other sources, including surveys, observations, and IT database records, then cross-correlating these with power data, the authors constructed a quantitative breakdown. This idea is discussed and further advanced in Chapter 4.

While standalone monitoring devices provide a straightforward mechanism to roll out energy measurement technologies today, they do not scale well to cover entire buildings; their configuration, deployment and maintenance becomes expensive. However, it is easy to imagine smart plugs and sockets becoming widespread in future, perhaps driven by a government requirement; this is a logical continuation of the current programme to roll out smart meters to every home. Every socket could contain circuitry to measure current and voltage similar to that embedded in the extension modules described above and report it, either wirelessly or using the power lines themselves for transmission. Plugs could even contain RFID tags or similar to allow smart sockets to detect the model of appliance connected; this could be looked up in an online database to obtain details and characteristics. This would allow maps to be built up automatically of all the devices in a building and their live energy usage. The problem of actual measurement is therefore not considered in great depth in this dissertation; real challenges lie in the analysis of this data.

2.3.1.4 Sub-appliance level

In order to divide the allocate the energy used by a shared appliance to the person responsible for causing the consumption it is necessary to identify the energy costs of specific actions it can perform. This problem is discussed further in Section 4.3.

Dutta et al. presented a simple design for energy metering *in situ* by augmenting switching regulators [48], and Fonseca et al. built on top of this hardware platform to apportion energy costs of components in embedded network devices to individual activities [59]. This allows developers to quantify the effects of different approaches, but requires significant hardware and operating system modification.

Flinn et al. have also contributed a significant body of work in the area of measuring and reducing the power consumption of larger mobile devices, including quantifying the energy consumption of a pocket computer [51] and the PowerScope tool for profiling energy usage [55]. These tools produced profiles of energy usage by process and procedure which could be used to reduce the consumption of adaptive applications. For sensor-based applications that run continuously in the background of the sort that may contribute important data to a personal energy meter, it is important to know how varying the frequency of measurements or how data is stored and transmitted might affect battery life; this problem is investigated further in Section 4.5.

Significant efforts have been made to reduce the energy consumption of wireless communication; while some system for measuring the power draw is required to evaluate these mechanisms, these have generally operated at a fairly coarse level. Pering et al. measured the voltage and current at the network interface cards, but sampled only every 10 ms and did not attempt to align the trace with specific actions [166]; similarly, Mohan et al. looked at the overall power required by the sensors for their pervasive application but did not investigate any further [145].

A new mechanism for decomposing power measurements of devices to determine the energy costs of relevant states and actions is required before the consumption of shared resources can be apportioned by a personal energy meter. Section 4.5 presents one suitable system.

2.3.2 Indirect electricity metering

Hart, Kern and Schweppe from MIT invented the concept of *Non-Intrusive appliance Load Monitoring*, or NILM, in the 80s with funding from the Electric Power Research Institute [83]. This uses building-level measurements, perhaps taken from outside by the utility company, and learning algorithms to infer which devices account for the load; it requires a training phase, then recognises step changes in the total load as appliances are switched on and off. Preliminary field tests showed that its results were usually within 10% of independent sensors provided when dealing with large appliances such as water heaters (2 kW) and dehumidifiers (700 W).

In a follow up paper, Hart also worried about the privacy implications of the technology he had invented [82]. There is significant potential for it to be abused, leading to an erosion of civil liberties; illicit printing presses could easily be discovered, and the sophisticated burglar could determine occupants' schedules in advance of his break-in.

In 1999, Drenker and Kader from the EPRI described progress with what they now called the *Non-Intrusive Appliance Load Monitoring System* (NIALMS); encouraging results in a beta test in 1998 led to offering the tested products for commercial sale [47]. They worked in conjunction with utility companies, each of which fitted several homes with both NIALMS and conventional device-level monitors. The study only investigated 'large' appliances as previously described, but the identification success rate was almost 100%

for those with only two states. For some appliances (heaters and pumps) the calculated total energy used was within 4% of the true value; for air conditioners, fridges and freezers it was around 13%. They claim that the commercial system has improved accuracy, and the technology is being extended to multistate loads.

NILM has a number of key advantages: only a single meter is required, minimising the time, cost and intrusiveness of the installation, and is a very promising technique. However, there are known shortcomings: uncertainty and undetected error, time and effort in calibration, and a restricted set of target appliances. It works remarkably well for large loads (over 150 W) that have few power states — either on or off, or with very simple operating states such as high, medium and low. Low powered loads and those with a large number of device states (like a dishwasher) or continuously variable energy usage (like an electric cooker) are very difficult to extract from whole-house measurements. Furthermore, the technique breaks down when applied to commercial facilities which may contain dozens or hundreds of indistinguishable devices. Most NILM systems also rely on processing data in batch using a day or more of stored data, making them unsuitable for real-time use. Laugman et al. explained and illustrated NILM techniques in more detail, extending them to noisy and commercial environments, and surveyed the systems developed over a period of 20 years as well [130].

Marchiori and Han tried to reach a compromise between building-level non-intrusive load monitoring and individual device-level metering by disaggregating measurements taken at circuit level. This has the advantages that there are fewer devices on each circuit and high-powered devices such as ovens tend to be installed on dedicated circuits so will not interfere with lower-powered ones; the trade-off is an increased hardware cost. They used an algorithm inspired by NILM but with probabilistic level-based, rather than edge-based, disaggregation. This makes it better suited to monitoring devices with complex state or continuously varying demands. Commercial meters were used, though any system (such as those described in Section 2.3.1.2) could provide the input data; provided some devices have a control system, and can therefore be turned off remotely, the system can perform automated training. In a trial with three devices (including a PC with variable power consumption) on a circuit the system achieved an average error after 24 hours of about 5%; however, in situ household trials have not yet been conducted.

The team behind the ACme monitors also tried using them to understand better where energy was being used in their building [105]. They used the concept of additive load trees, where each node represents the total consumption of its children; instead of metering every device directly, they evaluated the feasibility of reducing the number of sensors required by taking advantage of this additive property to infer the consumption of devices where the consumption of their siblings and parents are already known. While this is straightforward when only one node in a generation is not monitored, it becomes harder when there are several, and probabilistic techniques to distinguish between loads are tried, along with additional hardware such as light and vibration sensors. The authors also discussed extrapolating the consumption of a group of devices from a sample of the population and early attempts to disaggregate load spatially and by individuals as well as functionally. Their apportionment method was fairly simplistic, allocating consumption of owned devices directly to their owners but dividing the consumption of shared devices amongst everyone whose 'home coordinates' fall within the same enclosing space as the device. Nevertheless, they highlighted an important area for ongoing research.

ViridiScope combines data from magnetic, acoustic and light sensors monitoring signals emitted from appliances with measurements from a home's main power meter to learn and estimate device-level power consumption [116]. The authors used secondary indicators to infer a device's state rather than sensing its consumption directly, propose an automating sensor calibration framework and demonstrate its use in a two-bedroom apartment where it attained an accuracy of around 90%. The system can also support directly metered appliances, and groups together all uninstrumented devices as a single 'ghost appliance'. Although using indirect sensors eases the deployment task because they do not need to be installed inline with power cables, it is still necessary to deploy a sensor for each significant device in the building; furthermore, an extensive, albeit somewhat automated, calibration process is required.

Rowe et al. also attempted contactless sensing of appliance state transitions, using a method very similar to ViridiScope [182]. They found that the calibration burden to estimate power consumption of each device directly was too onerous, and instead focussed on detecting state transitions as an input for NILM to quantify consumption. This can help address some of the key challenges with NILM, namely the need for appliance-specific training and the problem identifying temporally-close transients. EMF detector sensor nodes as developed by the authors can be used in a continuous training process and to resolve ambiguities. They built on ViridiScope by moving filtering into the sensors instead of transferring raw data to a separate PC for processing.

Jung and Savvides considered the problem of disaggregating total power consumption by device based on on/off state sensing and propose a possible solution using load trees which can estimate its own accuracy [107]. They also proposed an algorithm for optimally placing additional power meters to increase the accuracy to a desired level. The method was evaluated in two case studies, with approximately 20% average prediction error.

Patel et al. investigated activity sensing based on detecting and classifying electrical events on a residential power line [164]. They relied on the electrical noise created by the abrupt switching of electrical devices and use machine learning techniques to recognise the patterns caused by individual appliances. Their focus was on understanding occupant activity, rather than power consumption, but the same data could be used as an input to any of the disaggregation methods described already; the technique only works for resistive and inductive loads and so will not help distinguish between electronic devices or others with switched mode power supplies. An evaluation in six separate homes found it achieved approximately 90% classification accuracy.

Some of the same authors subsequently introduced *ElectriSense*, which complemented the previous system by sensing the electromagnetic interference caused by switched mode power supplies [74]. This technique has the significant advantage that calibration can be performed once for a device and then used across homes, rather than requiring per-home calibration. Experimental trials in seven homes and one six-month deployment showed a classification accuracy of around 94%.

2.4 Metering other forms of consumption

2.4.1 Embodied energy

As discussed in Section 3.6, a personal energy meter could also try to account for the embodied energy involved in the manufacture, transport and ultimate recycling of the products we own and use. There is already a large body of work on Life Cycle Assessment which could provide the necessary data; this section gives an overview of a few representative examples of work in this field.

WattzOn[13] is a website that allows users to estimate their total energy footprint by answering a series of questionnaires with the stated goal of educating users about energy efficiency and conservation. It also features a embodied energy database containing details of the footprints of a significant number of consumer goods.[14] Users can select the items they own to have their costs added to their profiles. As with manual input of meter readings, this provides a straightforward way for users to derive useful results from a personal energy meter before widespread sensing can be deployed to automate the process.

Reichart and Hischier investigated the environmental impact of getting the news, observing that these assessments are not as straightforward as they might appear [173]. Using a single news item as a functional unit, newspapers appear to have a small environmental impact compared with news on television or web sites, but this only translates into reality if it is possible to buy parts of a newspaper. Considering the daily news as a whole, television fares significantly better than the Internet, which in turns is better than printed paper. Meanwhile, Shah et al. from Hewlett-Packard Laboratories argued that accurately quantifying information and communications technology's footprint is a critical first step toward reducing its environmental burden and provided illustrative first-order evaluations of both the embodied and operational costs of a handheld device, notebook, desktop, blade server and data centre [186].

Sourcemap[15] supports sustainable design through both life cycle assessment and supply chain transparency with shared visualisations [21]. It is a web-based tool that allows users to specify each ingredient or component of a product and where they come from; Sourcemap then estimates the embodied energy and draws a map linking each material at its source to a central assembly hub and thence to the location of the consumer. The user can also specify the lifetime energy consumption and planned end-of-life scenario to arrive at a whole-life-cycle assessment. These maps and estimates can be shared with the wider community, building a database of energy costs of common products; they can be created either by interested consumers or by the manufacturers themselves. The authors worked with five small business with a sustainability focus to help them create Sourcemaps of their products and highlight their environmental credentials through disclosure; they refined their own tool through the process. The resulting application is both easy-to-use and attractive, encouraging wider participation.

Dada et al. suggested that mobile phones provide a better interface to view estimates of the carbon footprints of physical products than static labels, since different instances of the same product may have footprints that vary spatially or temporally [36]. They demonstrated a simple prototype system using a phone that supported Near Field Communication, a short range technology compatible with existing passive RFID tags; products should be tagged with a unique ID which is used as a key in a unified database. Although RFID tags may be impractical, the use of mobile phones to interact with physical objects as a means to discovering their energy costs is adopted in Section 6.2.

[13] `http://www.wattzon.com/`
[14] `http://www.wattzon.com/stuff`
[15] `http://www.sourcemap.org/`

2.4.2 Transport

Energy used through transport accounts for about 35% of a typical individual's total consumption [136]. It is therefore valuable to understand how, and how far, each person travels, in order to estimate the energy cost of each trip.

The *Sentient Van* is a shared vehicle, available for use by all members of the research group that created it, which contains myriad sensors to record data automatically about each journey and its environment [39]. Its driver's identity can be inferred from the online booking system, or from the RFID transceiver mounted in the sun visor which can read ID cards. This latter system also allows the cost of journeys to be shared with a passenger whose presence would not otherwise be known. When the van is returned to its parking space it uploads details of its trips to a central computer for later analysis. With the proliferation of mobile satellite navigation devices and in-built sensors in cars it can be assumed that in the near future most vehicles will be capable of providing this level of data. However, it is also necessary to consider obtaining traces for other forms of transport.

Most new mobile phones include GPS receivers that might be used to record continuous position logs in the background. The first challenge is segmenting a personal GPS trace such as might be obtained from a smart phone in a pocket to identify individual journeys with start and end points and estimate the mode of transport used. Speed alone is enough to distinguish between walking, mechanised transport and aircraft; combining the trace with known map data such as the locations of train lines and bus stops helps distinguish public from private transport. Separating car journeys from cycling is often overlooked but can be difficult in congested city environments where speeds are likely to be comparable; additional sensors, such as accelerometers or even microphones, have proved valuable in this context. Zheng et al. have published significant work in this area [211, 212], while the Personal Environmental Impact Report project used the sensors in a mobile phone to determine if an individual is stationary, walking, running, biking, or in motorised transport and provided a brief overview and comparison of related work [172]. *UbiGreen* was a mobile phone-based application that provided personal awareness about green transportation behaviour through iconic feedback [61]. It attempted to infer mode of transport based on GSM signal, belt-mounted accelerometers and manual surveys. Thiagarajan et al. also used accelerometers and GPS data to infer when a mobile phone was being carried on a bus [193]. They observed that previously-suggested techniques, such as using a threshold on the user's GPS-derived velocity, are insufficient, since continuously running the GPS drains the battery, buses may travel at no more than walking speed in heavy traffic and GPS may not always be available. They therefore used the accelerometers to detect the difference between walking and vehicular transport (with 99.9% precision and 97.5% recall) before turning on the GPS and using adherence to the bus schedule, stopping at bus stops and inter-stop distance to distinguish buses from cars. Their system classified only 9% of bus trips wrongly, mostly in the cases of overlapping routes; this accuracy would be sufficient to provide valuable feedback to the user of a personal energy meter.

Air travel represents a significant proportion of an individual's total energy consumption. Most people take only a small number of flights each year and these tend to be planned in advance and well documented, making them easier to account for. Sites such as Dopplr[16]

[16] `http://www.dopplr.com/`

and TripIt[17] use natural language processing techniques to interpret the confirmation emails generated by airlines and travel agents and parse flight details. Furthermore, these sites encourage their users to submit details of forthcoming trips in order to benefit from their social networking features, pointing out where visits coincide with those of friends and suggesting hotels and restaurants others have recommended. Their growing popularity suggests that relying on users to volunteer information about their long-distance travel plans may be a simple and effective solution for the personal energy meter.

2.4.3 Water

A personal energy meter should be able to take account of all forms of energy usage, not just electricity or fossil fuels. According to the UN, water usage has grown twice as fast as the population during the past century. Today, already 1.1 billion people lack access to safe water. There is also a complex interplay between the production of electricity and water [117]—or example, the American public water supply and treatment facilities consume about 56 billion kWh per year, which is enough electricity to power over 5 million homes.[18] Accordingly, water consumption is an area of concern for personal energy metering. As with other utilities, smart meters are likely to become prevalent in the coming decades;[19] in the interim, a number of ubiquitous computing systems have been designed to fill the gap, and some researchers have started operating at a finer level to determine the consumption of individual appliances and outlets. These systems could provide valuable input data for a personal energy meter.

2.4.3.1 Building level

Monitoring water consumption on a fine-grained basis is traditionally quite difficult; flow sensors require installation within the pipe and so can generally only be installed by professional plumbers, and a large number of sensors would be required to separate the consumption of each device.

To get around this problem, Kim et al. proposed a *Nonintrusive Autonomous Water Monitoring System* (NAWMS), which used vibration sensors attached to individual water pipes to construct a self-calibrating system that provided information on when, where, and how much water they are using [115]. It relied on the existing water meter to provide details of overall consumption, but used the vibration sensors to estimate the flow to each individual outlet; since these sensors require calibration, the authors devised an adaptive self-calibration procedure. This, and the fact that the sensors were installed on, rather than in, the pipes, mean the system could be user-deployed. Unlike other non-intrusive systems it did provide data on actual consumption, rather than just user behaviour, and their testbed deployment showed a mean absolute error of 7%. Nevertheless, a commercial flow rate meter was required, which is expensive, and the process of deploying sensors on each pipe, specifying their topography and performing the calibration is time consuming.

Kim et al. also collected synchronised water and electricity usage data for 3 months in a family house using an ultrasonic water flow meter and a commercial mains power

[17] `http://www.tripit.com/`
[18] `http://epa.gov/watersense/index.htm`
[19] `http://www.utilimetrics.org/`

meter [117]. They observed that the combined traces reveal a lot of information about consumption patterns. Many water fixtures have recognisable usage patterns because they are mechanically or electro-mechanically controlled; their NAWMS technique detects this with good accuracy. However, smaller events such as sink usage and compound events remain difficult to detect; here they argue an electricity trace can help. Since in many cases both utilities will be metered anyway this seems a sensible approach to resolving ambiguities in either trace.

Fogarty et al. described a similar approach using microphones attached to the outside of water pipes to at critical locations to infer water usage and thereby activities within a home [56]. They focussed on activity recognition; the system could not infer water consumption, but could provide secondary indicators of usage which could be combined with profiles of water-consuming devices to come up with an overall estimate.

The *HydroSense* project also attempted to use the water infrastructure to determine human activity within a home, but in a significant step forward for incremental sensing used only a single pressure sensor screwed on at a water connection point [64]. It identified individual water fixtures within a home according to the unique pressure waves that propagate to the sensor when valves are opened or closed and, importantly for a personal energy meter, estimated the amount of water being used at a fixture based on the magnitude of the resulting pressure drop within the water infrastructure. The authors evaluated HydroSense in ten separate homes and found it could identify individual fixtures with 98% aggregate accuracy and estimate water usage with error rates "comparable to traditional utility-supplied meters"–though this was left undefined. However, to obtain this level of accuracy requires a careful calibration process, filling a bucket of known volume from each outlet in the house, increasing the complexity of deployment. Subsequently Campbell et al. created *WATTR*, a self-power water pressure sensor for HydroSense that eliminated the need for battery replacement imposed by the previous Bluetooth-based sensor [25]. This used the same pressure impulses as both a power and sensing source, and so further reduces the cost of deployment.

2.4.3.2 Device level

Bonanni et al. described a range of 'smart sinks', including *WaterBot*, which was a persuasive water conservation device that attached to a tap and provided context-sensitive feedback to its users [20]. It monitored the water flow and presented visual and auditory prompts by lighting up the stream, showing a continuous visual reminder of how long the tap had been on and providing 'positive auditory messages and chimes.' Users reported that it made them more aware of their water usage and more likely to turn off the tap when not required.

Kappel and Grechenig presented *show-me*, which adopted a similar approach of presenting feedback at the point of use but focusses on water conservation in the shower [110]. It consisted of a flow meter, microcontroller and feedback mechanism; a stick of blue LEDs would light up one-by-one as the shower was left on and so present an ambient display by way of a stylised representation of a column of water. In a user study with nine individuals in four households for three weeks the mean water consumption decreased by approximately 25%.

Both these devices relied on inline flow meters, which are unlikely in practice to be deployed to every water outlet; furthermore, it is more readily apparent to most people how they might reduce their water consumption and so the disaggregation is not as essential as for electricity. However, the projects do show that even very simplistic feedback mechanisms can have a positive effect on behaviour.

2.4.4 Gas

Natural gas is the most widely consumed energy source in many homes[20] and should therefore be included in a personal energy meter, but there has been little work in the field of ubiquitous computing attempting to understand where, why and how it is used.

Han et al. from LS Industrial Systems Co. of the Republic of Korea have built a wireless sensor network to read physical gas meters [80]. Unlike most systems described in the literature which are research prototypes deployed only in a handful of test homes, their units have been in use for actual billing purposes in 20,000 apartments in South Korea for three years. The system consists of sensing nodes, which attach physically to gas meters and use magnetic sensors to determine the meter reading, and relay nodes, which collate readings from hundreds of sensing nodes once a day and pass them on to the meter data management system. The sensors cost around US $20 each to manufacture, and have a battery life of at least five years. In terms of both hardware and installation labour this is an impressively low-cost alternative to replacing every single gas meter with a smart alternative.

GasSense is an example of infrastructure-mediated sensing which automatically monitors gas use and attempts to disaggregate it down to its source based on a single sensor [30]. A microphone is attached at the gas regulator, which controls the pressure of gas before it enters the meter. It does not require direct contact with the gas; this means it can safely be installed without the need for a professional. The resulting audio is analysed to infer the flow volume, which is shown to be directly proportional to the magnitude of the resonant frequency in the sample. A kNN classifier is trained for the appliances in each house and then recognises them based on both flow volume and rate-of-change; the aggregate accuracy across the sample of nine homes was 95%. However, these results were obtained in idealised conditions: the short duration of the trials means the effects of changes in temperature, humidity and pressure on the acoustic signal produced by the regulator were not accounted for; simultaneous events cannot be distinguished, and although the sensor can in theory be installed by the end user in this trial an extensive calibration process was performed in each house by the researchers. No results were given for accuracy of consumption estimates.

2.5 Reducing deployment costs

2.5.1 User deployed sensing

To be deployed on a global scale any additional sensing infrastructure required for personal energy metering must be inexpensive not only in monetary cost but also in terms of installation. Many systems, such as flow meters and circuit-level monitors, require professional deployment, which at a stroke makes them unlikely to be adopted in most

[20] http://tonto.eia.doe.gov/kids/energy.cfm?page=us_energy_homes

homes and small buildings. Sensors which can be installed without specialist knowledge are much more likely to succeed.

Beckmann, Consolvo and LaMarca explored the viability of domestic sensors being installed not by skilled researchers but by end users through an in situ evaluation in 15 homes of the installation kit for a hypothetical *Home Energy Tutor*, which is designed to help homeowners track their household energy use and learn about ways to reduce it [17]. The kit included vibration and electrical current sensors and microphones to detect the use of major appliances and motion detectors and cameras to monitor activity and electric lighting. Although many of the sensors included were non-functional mockups, they are representative of the kind of devices that would be required for an application not dissimilar to a personal energy meter and so the study is of particular relevance. Users were required to place sensors on appliances and in rooms around their homes and create an association between sensors at their targets by scanning their barcodes with a handheld computer; different types of sensors had different requirements in terms of placement (both position and directionality) and association. Overall, 115 of the 150 sensor installation tasks were completed successfully; 5 participants completed all the tasks, while 2 failed to complete any. The failures reveal more about principles to which systems should adhere so as to maximise the chances of acceptance and straightforward end-user installation:

1. Make appropriate use of user conceptual models for familiar technologies
2. Balance installation usability with domestic concerns
3. Avoid use of cameras, microphones, and highly directional sensors if possible
4. Detect incorrect installation of components and provide value for partial installations
5. Educate the user about data collection, storage, and transmission

They also found many negative reactions to the intrusion of sensors into the living space, including objections to the potential for damage caused by the adhesive used for installation, concerns that sensors were placed in locations accessible by children or pets, and objections to the placement of cameras and microphones in the home.

Kawsar et al. also investigated do-it-yourself deployment of ubiquitous computing systems in a home context and presented an infrastructure that provided the foundation for involving end users in the deployment process through manipulating RFID cards [111]. Although much of the framework revolved around assembling control applications and required tasks to be written in a specialised manner, their custom-built end user deployment tool which allowed end users to install, uninstall, run and stop artefacts and applications was a lower-cost alternative to the handheld computer used by Beckmann et al. and a user trial was a success. 25 participants were asked to attempt 4 deployment tasks each; all of them successfully finished the assigned tasks, though 6 needed active support; subjective evaluations were positive. The system introduced significant overheads, with dedicated additional hardware; this is likely to prove an obstacle to adoption in itself, but their experiment confirmed the principles listed above to which sensors should adhere.

2.5.2 Crowd-sourcing inventory information

Constructing an inventory of all the energy-consuming devices in a building and their owners is laborious and error-prone, and keeping it up-to-date requires ongoing effort; audits carried out by a single person or team tend to diverge from reality shortly afterwards as occupants acquire or move appliances without notifying building services.

Fortunately, a inventory need not necessarily be perfect to be useful for estimating energy consumption. In 2010, the vast majority of energy consumers in the William Gates Building were low power electrical devices consuming less than 100 W. Individually, each device accounts for less than 0.05% of the total building consumption, although their combined consumption does amount to a significant figure. The inventory therefore need not contain every single device as long as it contains a *representative* sample and an estimate of its coverage (devices missed that are unusually power-hungry will skew the result). Obviously, confidence in the sample will increase with its size.

The phenomenal success of Wikipedia is a prominent example of the ability of a large community to generate and maintain an enormous repository of knowledge, motivated not by necessity or reward but by factors such as fun and ideology [152]. The OpenStreetMap project has demonstrated how volunteers can build accurate and detailed maps comparable with commercial counterparts simply through the annotation of data collected through commodity GPS receivers [78]. OpenRoomMap explores how such techniques

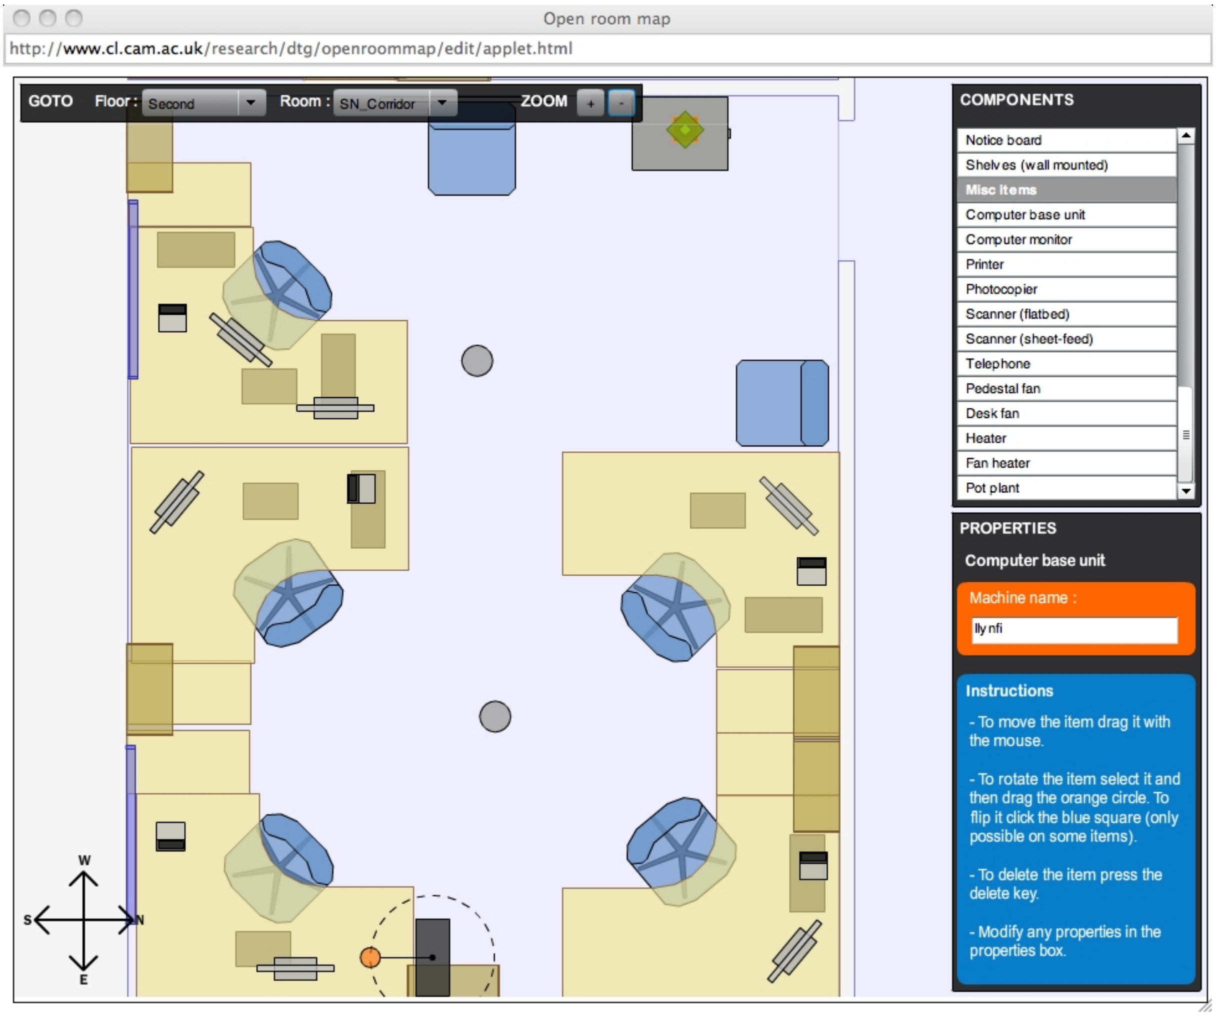

Figure 2.2: The OpenRoomMap interface provides an editable plan of the building and a toolbox to add new items [178]

could be applied to make creating and updating the inventories that are required easier, more accurate and more enjoyable.

Although one-off surveys by designated individuals are clearly a possibility for personal energy metering, a possible alternative is OpenRoomMap, a crowd-maintained building inventory system initially developed by Woodman and Rice at the University of Cambridge which delegates the job of maintaining an accurate building inventory to the occupants of the building [178]. This section reviews the system and describes some extensions and new results on its uptake after it was deployed.

An OpenRoomMap model consists of a non-editable floor plan on which instances of defined object types can be dragged (Figure 2.2). Object types currently include furniture as well as electronic equipment including computers, monitors, printers and telephones. Certain object types also have semantic data fields that can be edited by the user—for example, each instance of a telephone has an associated telephone number and each computer has a name, which allows network traffic to be used as a secondary indicator for power state. For the sake of simplicity, the current model does not encode relationships between objects. For example, a monitor is not explicitly associated with the computer to which it is attached. Nevertheless it is often possible to automatically infer such

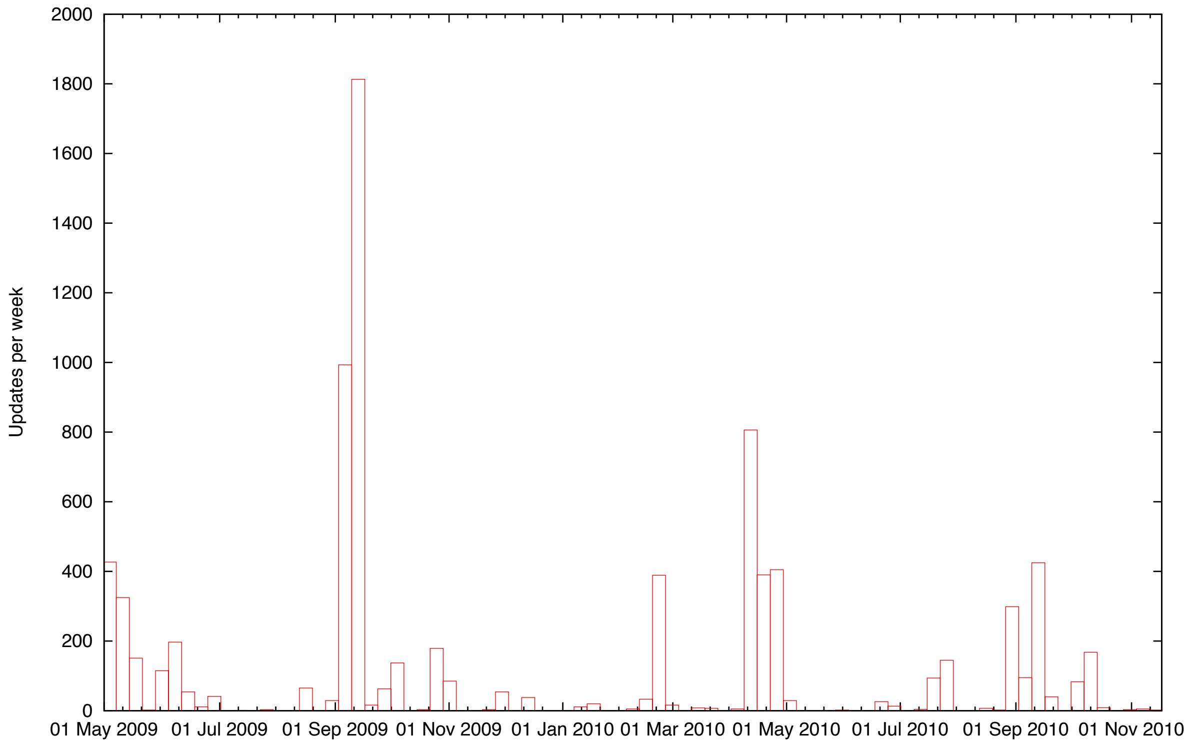

Figure 2.3: Number of updates made to OpenRoomMap each week since launch.

relationships from the model, for instance by considering the distance between two objects.

The crowd-sourcing aspect of the data collection means it is much more likely that a representative selection of the building inventory will be collected than by attempting to persuade building managers to perform surveys manually. As an example, OpenRoomMap was launched by email to staff members at the Computer Laboratory who placed around 3,000 objects over 5 hours. After a period of approximately 6 months the map was approximately 70% full with 139 out of 207 rooms completed [178]. To determine whether this usage was sustained, further investigation was carried out. By December 2010, 183 of the building's 313 users had updated the map; 104 of these had made more than 50 changes. 5,988 separate items had been created in total, and 1,427 deleted. Figure 2.3 shows the number of updates made each day after launch (the first day had over 7,000 updates and so is not included). The clear peaks correspond to the start of each new term, when new students join the department and offices and equipment are therefore reassigned. Although the crowd-sourced data is not perfect, it is the most reliable inventory available for the building; production use of the system to display printer details, book meeting rooms and help visitors locate offices helps ensure it is kept up to date since the same people who can update it benefit directly from its accuracy.

An informal email survey of 15 of the top contributors revealed four main reasons why they chose to update the map, listed here under the same motivational categories proposed by Clary et al. [29] and also used for a study of what motivates Wikipedians [152]:

Understanding 7 people contributed to the map because they themselves found the data useful. One specifically mentioned its potential value for monitoring energy consumption, although this was not part of the original release.

Fun 6 people used it simply because it was fun or provided a good excuse for procrastination. One commented "I enjoy getting it to match reality;" another "I like world-building games like Civilisation, and I think the appeal is the same."

Values 4 people were motivated by altruistic values, stating "I admire user-edited projects like Wikipedia" and "I wanted the project to succeed."

Social 3 people cited social reasons, including peer pressure: "I also remember noticing that some people in my corridor already had their rooms mapped while mine wasn't."

Additionally, 2 people commented that they were interested in finding out more about the technical aspects, features, limitations and interface of the tool; this is less likely to be generally applicable outside the research community.

To better fit the needs of apportionment for personal energy metering, OpenRoomMap was extended to store details of the owner of each item in its database. Each room or space can be associated with one or more occupants. To estimate ownership from the existing world model in OpenRoomMap, it is assumed that devices in a room with only one occupant belong to that person, while those in a room with more than one occupant are owned by the person who placed them on the map, provided he is is one of the occupants. This allows estimation of a personal load for each individual based on the devices he owns; in office environments there is often a relatively homogeneous set of equipment, so a device can be profiled once and that profile used for each instance of the device.

2.6 Occupancy detection

Occupancy information alone is useful in refining initial estimates of energy consumption (Section 3.5). In the field of detecting a user's presence, Garg and Bansal showed how to improve on the estimates of simple occupancy sensors by adapting to changing activity levels [65]. Dodier et al. explored the use of belief networks with occupancy sensors [45]. There is a significant body of work on simulating occupancy profiles using Markov chains where live data is unavailable [179, 196].

To determine *who* is in a space at any given time or keep track of the movements of named individuals requires a system more sophisticated than the occupancy-detection mechanisms surveyed in this section. This more fine-grained data is necessary to divide up the consumption of shared resources accurately (Section 5.1), and the following sections set out the key requirements for a location system for personal energy metering and provide a thorough survey evaluating existing systems against these requirements. Chapter 5 demonstrates a novel mechanism for tracking location with minimal additional infrastructure.

2.6.1 Force-based

A person standing on the ground exerts a force on it equal to his weight; at the same time, the ground exerts an equal and opposite reaction force on the person. By instrumenting the floor to measure the application of this force, it is possible to determine the locations of people and objects on it; additional knowledge about possible weight signatures may allow their identities to be inferred. These so-called “active floors” appeal because they promise low-power, inexpensive sensing with minimal hardware deployment and no user augmentation.

The first example was the *Active Floor* from Olivetti Research [4]. This consists of a square grid of conventional 0.5×0.5 m carpet tiles, supported at the corners by load cells which are instrumented to give the total vertical force. Bending moment analysis is used to determine the centre of pressure, which is equated to the position of a user or object. Clearly, the system cannot distinguish between two people standing on the same tile, so smaller tiles increase reliability and resolution but also deployment costs and effort. ORL used Hidden Markov Models to attempt to identify users from a set of 15 known subjects; this achieved some success, but clearly limits the utility of the system to a small number of regular users rather than occasional visitors.

As part of the *Aware Home* project to build technologically advanced house to serve as a living laboratory, the Georgia Institute of Technology built a very similar hardware system [154]. They used a nearest-neighbour search over a feature space to match footsteps against a set of known users, and achieved a 93% overall user recognition rate from a set of 15 subjects.

As an indication of ease of deployment, the authors wrote that they intended to install ten Smart Floor tiles in strategic locations in the house, including house entrances, hallway entrances, kitchens, and bedroom entrances; cost prohibited augmenting the entire floor.

All these systems rely on expensive components and deployment is complicated. Subsequent attempts to address the problem have instead used a mesh of wires. For example, the *Magic Carpet* developed at MIT uses a grid of piezoelectric wires hidden under a 6×10 foot carpet to monitor foot position and pressure [159]. The *POISE* system (POsition Information through Static Electricity) can localise foot falls and lifts with a floor augmented to sense local electrostatic charges [49]. This augmentation amounts to a series of standard, low-cost, PVC-sheathed wires. The wires draw no power, acting as passive sensors. This approach has the advantages of being extremely low-cost (a few pounds for the prototype pad), drawing minimal power (milliwatts) and being straightforward to deploy and configure. Unfortunately, the passive sensing is unable to distinguish between individuals without other sensor data; with no pressure sensors it cannot identify users through their known weights.

2.6.2 Sound-based

Tarzia et al. proposed a technique to detect the presence of computer users via ultrasound using only hardware that already exists on commodity laptops [192]. The system works by detecting the changing reflections of a generated sine wave caused by a movement. A study with 20 volunteers found it was possible to detect the presence or absence of users with near perfect accuracy after around 10 s of measurement. The ability to repurpose existing infrastructure in this way makes it a much more realistic candidate for widespread deployment, since no additional deployment costs are involved. Unfortunately, as with active floors and other tagless systems, it is unable to differentiate between individuals; it can only infer presence, not identity. However, this may well be an appropriate technique for learning when a user is in his office for the purposes of apportionment if he has a personal computer that is not shared with anyone else.

2.7 Purpose-built location systems

This section provides a review of purpose-built indoor location systems described in the literature and assesses their suitability for use in personal energy metering. In general, it is difficult to compare different systems since their performance is dependent on the physical parameters including size, furniture, layout of walls and partitions and beacon positions. Nevertheless, several surveys have been published on this subject [71, 93, 94]. Unlike most surveys, this section does not focus on features, accuracy or resolution, but rather on those properties such as cost and ease of deployment which have been identified as crucial for the widespread adoption necessary for a personal energy meter. It only includes technologies that have the potential to track users continuously, and therefore excludes techniques such as barcode-based [174] or markerless [125] visual tracking which are designed for user-driven, interactive localisation. Although these could in principle be implemented on today's mobile phones, they require that the camera always has a good view of its surroundings, which is rarely the case.

2.7.1 Infrared

Infrared signals can be used to transmit information wirelessly and also to provide location information: since they cannot pass through walls, if a signal transmitted from one device is received by another then it is highly likely that they are located in the same room. Infrared light is strongly reflected by walls, meaning that line-of-sight is not required, and this mechanism also has the advantages that infrared transmitters and sensors are cheap and require little power. These characteristics make it a good candidate for room-level location.

2.7.1.1 Active Badge

Olivetti Research pioneered the concept of indoor location systems with their development of the archetypal *Active Badge* in 1990 [198]. These badges were worn by employees and would emit a unique code at 15 s intervals via pulse-width modulated infrared, which was picked up by a network of sensors placed around the host building. A master station polled the sensors for badge sightings, processed the data and made it available to clients.

A badge is a culturally-accepted identification and access control system, but to be viable it must be small, light and convenient to wear. To boost user acceptance, the Active Badges were a designed in a convenient form factor to wear on clothing ($55 \times 55 \times 7$ mm), and weighed around 40 g; the infrequent signals and a light sensor that turned the badge off when it was dark meant batteries lasted about a year.

This early system has in fact been more widely deployed and adopted than most research prototypes: at the time of the original paper there were over 100 badges, 200 sensors and 5 badge networks in use at 4 sites in Cambridge. A follow-up paper introducing an authenticated variant explains that a year after the introduction of the system and with no requirement for compliance all employees continued to wear badges, and that together with an installation at the University the user group size was about 150 people [197]. By 1994, when the badge evolved to use a hybrid RF/infrared approach in an effort to make the system more scalable and offer desk-level location, over 1,500 badges and 2,000 sensors were deployed throughout the research community in Cambridge and at a number of other universities and laboratories[21] in Europe and the USA [84].

These are impressive statistics, but they nevertheless highlight a problem with this approach: a large number of sensors are required, and although these can be manufactured relatively cheaply their installation is a significant undertaking. They need to be placed high up on walls or ceiling tiles of offices and on the entrances and exits of corridors and other public areas, and they must all have a power supply and be connected to a common, wired network. The original application for the Badges—routing telephone calls to wherever the intended recipient was located—has been rendered largely irrelevant by the advent of mobile phones, and the burden of deployment is probably too high for most outside the research community to consider.

2.7.1.2 ParcTab

The ParcTab system developed at Xerox PARC is widely held as one of the first examples of ubiquitous computing [2, 200], and its history was closely tied to the parallel developments taking place at ORL [183]. It consisted of palm-sized mobile computers that can communicate wirelessly through infrared transceivers to workstation-based applications. It adopted a very similar mechanism to the Active Badge, transmitting a beacon every 30 s even when idling in low-power mode that allowed the system to continue to monitor each ParcTab's location.

The first ParcTab system released in March 1993 consisted of 20 users and 25 transceivers. By the time of the second release in 1994 there was community of about 41 users and 50 transceivers at the Xerox PARC Computer Science Laboratory, and a further 10 transceivers at EuroPARC.[22] Again, the cost and effort of deploying the necessary infrastructure seems to have hindered wider adoption.

2.7.1.3 Smart Badge

A further derivative of the Active Badge concept was the *Smart Badge* [16], which also used infrared for location in the same manner but added a collection of sensors and actuators, allowing it to discover and control its physical environment and report data using the infrared communication channel. Once again, the project seemed never to get past the prototype stage, with the authors reporting that just ten Smart Badges and ten Badge Transceivers were constructed.

2.7.1.4 Locust Swarm

The *Locust Swarm* improved deployability by inverting the Active Badge methodology and adopting an inside-out approach: instead of individual badges broadcasting their

[21] `http://www.cl.cam.ac.uk/research/dtg/attarchive/ab.html`
[22] `http://www.ubiq.com/parctab/`

identifiers and being detecting by a network of sensors, a 'swarm' of independent Locusts were deployed throughout the building which broadcast their location to interested devices [118]. This removed the need for Locusts to be interconnected, and the clever inclusion of a solar cell meant no separate power source was required either. Locusts were instead placed in the grilles beneath overhead fluorescent lights, and covered a 20 ft diameter circle underneath them; no batteries or mains connection were necessary, easing installation and maintenance. In 1997, 300 units were being deployed throughout the MIT Media Lab after a successful trial with 10 systems. However, a large number of bespoke sensors must still be installed in carefully surveyed locations.

2.7.2 Ultrasound

2.7.2.1 The Bat System

The concept of using ultrasound for indoor location and context-awareness was first introduced by Ward, Jones and Hopper from the University of Cambridge and the Olivetti and Oracle Research Laboratory (ORL) as a follow-up to their successful experiments with the Active Badge [201]. Measurements are made of times-of-flight of sound pulses from an ultrasonic transmitter to receivers placed at known positions around it. Transmitter-receiver distances can be calculated from the pulse transit times, from which, in turn, the transmitter's location is found by multilateration. They expanded on the principles and described prototype hardware; with 16 ceiling receivers operating over a volume of 75 m^3, 95% of raw readings lay within 14 cm of the true position. This prototype developed into the *Bat system* [3, 85].

The system required deployment of not only a dense array of network ultrasound receivers but also radio transmitters and thermometers in each room (used to compensate for the effect of temperature on the speed of sound). The location of each ceiling-mounted receiver had to be surveyed as accurately as possible; even millimetres of error degrade localisation accuracy significantly. AT&T Labs Cambridge developed a mechanical measurement system called the crate for this purpose, which employed three one-dimensional measurement sensors[23]—retractable steel cables whose current extended length is measured with shaft encoders. The measurement sensors were mounted at known locations on a large rigid metal frame to triangulate the position of objects in three-dimensional space. By touching the end of each cable to a survey point, trilateration could be used to determine the points' locations. When the system was subsequently installed at the University of Cambridge theodolites were used to survey the deployment. Both of these devices are expensive and cumbersome, and Scott and Hazas report that with either method a typical room took one person-hour to survey [185].

The Bat system is probably the most accurate of all indoor location systems, but it has only ever been deployed in one or two locations at a time due to the prohibitive costs in both time and money involved in installing all the necessary equipment. It is therefore a valuable research tool but not a practical, general solution for personal energy metering.

[23] `http://www.asm-sensor.de/`

2.7.2.2 Cricket

The *Cricket Location Support System* used a very similar principle to the Bat system but was designed to address privacy concerns arising from its centralised design and therefore adopts an inside-out approach similar to that of GPS [169]. Instead of deploying a network of receivers, a series of beacons were installed at known locations. They periodically both emited an ultrasonic pulse and broadcast their location over a radio channel, allowing mobile devices which received both signals to infer their own positions independently. The beacons were uncoordinated, using independent, randomised transmission schedules; although this degrades positioning accuracy it does make their deployment simpler and reduces overall system costs.

As described in the original paper, the system could provide a location granularity of 4×4 ft by placing the beacons in a 4×4 ft grid. A subsequent modification to the system, the *Cricket Compass* improved on this to provide not just position but also orientation, and increased the resolution to centimetre-level, at the cost of a more dense beacon deployment [170]. Nevertheless, almost all of the disadvantages of the Bat system also apply to Cricket.

2.7.2.3 Dolphin

The *Dolphin* system by Hazas and Ward was an evolution of the Bat system. It replaced the standard narrow-band ultrasonic transceivers with custom ones capable of generating broadband signals, which are more resilient to noise and interference. This helped avoid some of the problems exhibited by the Bat system, which would sometimes accurately determine the position of a crisp packet being opened, or keys being jangled, instead of the Bat it was seeking. Broadband signals also allow multiple devices to be addressed simultaneously, meaning the update rate does not drop as the number of devices increases.

Several variations on the scheme were explored in a prototype deployment, although no system suitable for production was ever constructed. The technique can be used to build either a centralised system with an architecture similar to that of the Bat [91] or an inside-out, privacy-preserving variant [92]. In either case, the positioning accuracy was similar to the Bat system (better than 5 cm in 95% of cases). The Bat system required all the ceiling-mounted infrastructure to be connected to a common network. With the Dolphin system it was possible to reduce the time required for installation by using battery-powered transmitters and radio-based synchronisation, meaning only placement and surveying is needed; the tradeoff is an increased regular maintenance burden.

The system suffered from most of the same barriers to widespread deployment as the Bat does—extensive arrays of sensors still had to be installed in the ceilings in carefully-surveyed locations, and in the inside-out system their locations must be distributed to all the clients to allow them to position themselves. Furthermore, the broadband hardware is more complicated and expensive, relying on custom-built components, and the positioning algorithm depends on Fast Fourier Transforms to correlate the received signals with the expected spread spectrum waveforms. This is computationally very expensive—the operations were run on a workstation PC in the prototype described—and the authors noted that it is likely that specialised hardware correlators would be needed in a deployable version of the mobile receiver unit.

2.7.2.4 Other

Constellation was a combined inside-out ultrasound and inertial system designed for tracking both position and orientation with the high accuracy and update rate required for augmented reality applications [60]. Data from triaxial accelerometers is double-integrated to provide high-frequency position updates, while time-of-flight based ultrasonic range measurements correct the drift that is incurred. This provides tracking accuracy of the order of millimetres at a rate of around 500 Hz—comparable to optical systems—but at an extremely high hardware and deployment cost. The authors mentioned that accurately measuring the beacon locations can be time consuming and a dominant source of tracking error, and suggested an auto-calibration mechanism might be possible in future.

Randell and Muller described a simple combined RF/ultrasound system which they claimed could be implemented for around US $150 [171]. This system did not aim to rival commercial systems, but they claimed accuracies of 10–25 cm. A short range FM transmission was used to synchronise the mobile device with the four fixed ultrasound transmitters mounted in the ceiling. The transmitters then emitted short chirps, which are detected by a receiver with a PIC microcontroller that is either connected to a hand-held device or worn on the shoulder. The PIC is programmed to measure the number of 100 μs delay units occurring between transmission and reception of each chirp. These delay units correspond to 3.4 cm, giving an optimum resolution of 2.4 cm at one and a half metres below the centre of the transmitter square. Although it is much less expensive, it still suffers from the same problems as other ultrasound systems, requiring dedicated hardware to be installed in carefully-surveyed positions in the ceiling. A further major limitation was that it did not provide continuous resolution between cells—that is, the system was limited to four receivers to cover a single room, which is sufficient for research purposes but means a user cannot be tracked as he moves throughout a building.

2.7.3 Radio-based

Ubisense is a commercial indoor location system developed by many of the team behind the Bat system which uses ultra-wideband radio to position tags to within 13 cm of their true positions 95% of the time.[24] The wide frequency range of UWB pulses mean they can be extremely short, making reflected signals easier to detect. The main benefit of radio over ultrasound is that the signals can pass through many materials, meaning line-of-sight is no longer required between tags and receivers and so the number of receivers required is far lower; unfortunately, although the system's properties make it well suited for energy metering, the infrastructure remains very expensive and so it is unlikely to be deployed for that purpose.

SmartMoveX was an RF-based active badge system from Microsoft Research designed to be low-cost [126]. Users wore small transmitters; receivers around the building reported signal strength readings to a central PC which attempted to match them against a library of training vectors to determine the user's location. They estimated the parts for a badge cost $6 and for a receiver $16; the receivers connected to existing PCs on users' desk and so used the existing network infrastructure, further lowering costs. Four receivers were used to cover an area of 350 m^2, giving a cost excluding badges of around $0.18 per

[24] `http://www.ubisense.net/`

square metre. This is a significant improvement on most other active badge systems, and the concept of reusing existing PCs as base stations is promising and developed further in Section 5.1.1. However, gathering the training data requires a very significant time commitment, discussed in Section 2.8.3.3; if this is not considered prohibitive, as it will be in many environments, the same technique could be used with existing radio transmitters such as WiFi, DECT or FM stations, reducing the infrastructure cost further. Systems based on these ideas are surveyed in Section 2.8.

2.7.4 Inertial

Inertial navigation is a promising approach for infrastructure-free indoor tracking. Initially developed for aircraft and submarines where opportunities for position fixes are few and far between, advances in Micro-Electro-Mechanical Systems (MEMS) technology means accelerometers and gyroscopes are now sufficiently small and cheap that they can be attached to individuals as well. These devices suffer from significant inaccuracies, and it has been demonstrated that general-purpose inertial navigation algorithms are unsuitable for pedestrian tracking due to the rapid accumulation of errors in the tracked position. However, purpose-built systems have produced impressive results: by detecting when the foot is stationary and applying zero velocity corrections a pedestrian's relative movements can be tracked far more accurately than is possible using uncorrected inertial navigation.

Woodman has done pioneering work in this area [206]. He used particle filters to enforce constraints such as impassable walls and floors and thereby narrow down the absolute position of a pedestrian as he moves through an indoor environment [207]. Once the user's position has been uniquely determined the same filter can track this absolute position to sub-metre accuracy.

Regrettably, the Xsens hardware used cost in the region of £1000,[25] and cheaper devices have not been shown to provide the level of accuracy required for position to converge. Furthermore, the inertial tracking unit must be worn on the foot, which is limiting in a real-world situation; although Nike have demonstrated an accelerometer-based device for running training which fits neatly into a training shoe and transmits data wirelessly to an iPod[26], in general it is unlikely that such hardware will be integrated into every shoe a user might wear.

Modern smartphones typically have triaxial accelerometers built in to support gesture recognition and novel interfaces; the iPhone 4 was the first also to feature a triaxial gyroscope,[27] and it seems reasonable to suppose that other major vendors will follow suit. The ideal would be to make use of these sensors to perform inertial navigation continuously in the background when position fixes from GPS or other sensor systems are not available. At present the error they introduce is too great, but this is likely to improve as MEMS technology matures; the most significant challenge is that phones are not generally carried in a known location and not even fixed rigidly to the body. This makes even step detection difficult, and prevents techniques such as the footfall zero-velocity update relied upon by current state-of-the-art tracking systems.

[25] `http://www.xsens.com/en/general/mti`

[26] `http://www.apple.com/ipod/nike/`

[27] `http://www.eetimes.com/electronics-news/4199985/First-MEMS-gyro-packed-into-smarter-iPhone-4`

Through observation of 419 urban residents, Ichikawa et al. found out that most (around 60%) men carried their phones in a front trouser pocket while women carried theirs in a handbag [102]. Even sensing the device's own location is difficult using accurate accelerometers alone at present; Jin and Fujinami achieved an average of 73% accurate online classification between five different on-body locations [106], while Kunze et al. managed success rates in the low nineties while the user is walking [127].

As a small step towards the larger goal of inertial navigation, Blanke and Schiele attempted to recognise known location trajectories within buildings using a Xsens inertial navigation unit with unknown placement and orientation in a trouser pocket, achieving classification rates of about 95% on average [18]. Steinhoff and Schiele then conducted an experimental study of several approaches for dead reckoning in this scenario with unconstrained placement of a device in the user's trouser pocket; because the devices position inside the pocket is quite stable and the motion and rotation are coupled with the thigh motions they believe this location is promising [189].

Inertial navigation using dedicated foot-mounted accelerometers is clearly impractical for most users of a personal energy meter, but it remains an interesting avenue of research with much potential, particularly as the sensors embedded in phones improve and new methods are discovered to tackle the problem of unconstrained placement. Another limitation is the drain on the phone's battery from the continuous sensing required.

2.8 Opportunistic location systems

Another approach to reducing installation costs is to leverage existing infrastructure that has been installed in the building for a different purpose and exploit its characteristics to obtain contextual information. This section surveys a number of systems that adopt this technique.

2.8.1 Sound

Beep was a 3D location system which used audible sound for positioning [138]. This has the significant advantage of being almost universally supported by existing mobile phones and other devices, meaning no additional infrastructure at the user end is required. It used a similar technique to the Bat system (Section 2.7.2.1): when a user requested positioning, his roaming device synchronised with dedicated acoustic sensors though a wireless network and transmitted a pre-defined audible signal. The sensors detected this signal using specialised digital flights and made an estimate of the time-of-flight, which was reported to a central server. This performed multilateration to derive a position estimate and reported the results to the roaming device. Heterogeneous mobile devices introduced unknown hardware delays between a positioning request and the emitted signal caused by the sound card. The authors described a novel multilateration algorithm that takes this into account, and studied the number of sensors required to cover a given tracking volume. They concluded that to cover a warehouse of area 10,000 sq. ft and height 10 ft with sensors having range 25 ft would require between 30 and 38 sensors, depending on the deployment pattern. The testbed used six desktops with microphones placed on the ceiling as the sensors in a room measuring $32 \times 18 \times 8$ ft; their results showed that in a

2D plane the system has an accuracy of about 2 ft in more than 97% of cases, while in the 3D case this dropped to about 3 ft in 95% of cases.

The system as described was dependent on mobile devices requesting they be localised and the central server queueing these requests to avoid interference and is therefore more suitable for interactive applications than ongoing, passive tracking, though the same underlying technique could be used to achieve this.

An obvious downside is that the signals themselves are audible to humans; although their duration is kept to 100 ms to minimise irritation, the more devices being tracked and the more frequent the updates the more distracting this will become. Noisy environments also degrade the reliability of the system. The testbed used dedicated desktop PCs as sensors; while this is clearly a barrier to deployment, the authors proposed to use custom hardware in a future development, which they argued could be built at low cost and require little power. An alternative approach might be to use the desktop PCs already in use throughout office buildings as sensors, though their locations, typically on the floor under desks, are unlikely to be well-suited to multilateration applications.

In a similar vein, *WALRUS* (Wireless Acoustic Location with Room-level resolution using UltraSound) is a system that uses the built-in wireless networking interfaces and microphones of mobile devices to determine location with no special hardware [22]. It has a similar design to the Cricket system, with unsynchronised beacons periodically transmitting both their location via radio and a pulse of ultrasound, but uses PCs with standard desktop speakers to transmit ultrasound and WiFi as the communication channel instead of specialised radio.

Client devices do not attempt to perform multilateration; instead they record audio whenever they hear a location broadcast over WiFi, and if subsequent analysis finds energy in the ultrasonic range it is likely the device is in the same room as the beacon that generated the last packet since ultrasound tends to be blocked by walls. Tests showed the system is robust to background noise.

The lack of synchronisation between beacons means the chances of collisions increases with the number of rooms. With only 6 beacons, the mobile client was able to determine its location correctly about 84% of the time, but this dropped below 50% with 50 beacons. This may limit its viability in larger buildings—although a centralised, round-robin approach like that used by the Bat system could be adopted. The use of 802.11 broadcast packets does preclude clients from using a wireless network and the location system at the same time; addressing this would require changes to both access points and network card device drivers. These two constraints combine to make the system impractical for continuous tracking of users for a personal energy meter, but its philosophy of repurposing the PCs on every desk as base stations is a good approach that is adopted in Section 5.1.1.

2.8.2 Home infrastructure

A team at the Georgia Institute of Technology introduced *PowerLine Positioning*, or PLP, which used the residential powerline as the signalling infrastructure to create a whole-house indoor location system with sub-room-level accuracy [165]. They called this approach ‘infrastructure mediated sensing’, or ‘home bus snooping’ (as opposed to ‘direct sensing’). It required only the installation of two small plug-in modules at the extreme ends of the home and claimed an accuracy of 87–95% for classifying regions at 3 m and 67% at 1 m resolution; however, subsequent repetitions achieved only 54% room-level accuracy [190]. On closer inspection, there are a number of serious problems that limit its practicality. Most importantly, like all fingerprinting solutions, it is susceptible to accuracy degradation over time due to variability in the fingerprint. This is particularly problematic in this system since it requires a pair of frequencies independently injected into the power lines at separate points in the house. It is not possible to select a specific pair of frequencies that can be guaranteed to work in every setting over a period of time, necessitating frequent labour-intensive recalibration of the entire system.

In a subsequent paper, Stuntebeck and other members of the original team proposed the use of wideband signals, which are less susceptible to temporal variation, to counter this problem [190]. This also has the advantages of improving accuracy by providing additional features to the classifier and working in commercial spaces. The system they described required a trolley full of equipment, making it impractical for real-world deployment, but the authors argued that their prototype has more functionality than is necessary and described how a deployable version might be built.

The same team from the Georgia Institute of Technology have explored the possibility of detecting human movement and room transitions by using a single sensor in the HVAC system ductwork [163]. Disruptions in airflow, caused by human movement including room-to-room transitions and doors being opened or closed, result in static pressure changes in the HVAC air handler unit. Sensors mounted on the air filter allowed them to detect these pressure variations and classify where in the house the events are taking place with up to 75-80% accuracy. Occupancy of individual rooms could be derived probabilistically from a series of room transition and door events.

Clearly this approach suffers from the same problems as all occupancy and movement detection systems in that it is impossible to identify individuals without additional sources of information, but given this constraint the installation of a single instrumented air filter is clearly far simpler than that of an array of motion detectors throughout the house, each of which must be carefully aligned and surveyed.

One problem with relying on home infrastructure is that the techniques are very dependent on the style of building construction, and therefore in many cases their applicability is geographically limited. Powerline positioning does not work in most offices and is unlikely to work without significant modification in homes in Europe since they tend to be of a very different construction to those in Georgia; clearly HVAC-based systems are not an option in countries whose climates do not require air conditioning. They may yet be a useful tool in the arsenal, reducing the need for additional sensors in situations where they can be applied.

2.8.3 Radio techniques

Several techniques have been devised to take advantage of radio beacons that are already widely deployed. Individual systems are described in the following sections, classified by technology; this section provides an overview of the possible mechanisms which can be applied regardless of the radio source used.

2.8.3.1 Proximity

Proximity-based systems are the simplest conceptually. They assume that if a target is in range of a fixed base station then it must be approximately co-located with that base; this can be extended to assume that the target is co-located with the base from which it receives the strongest signal. Their resolution is therefore determined by the radio range, making this technique well-suited only to relatively short-range technologies such as Bluetooth if, as for the personal energy meter, indoor, approximate room-level location is required. Such systems cannot provide coordinate locations for the targets, but the major advantage is that no information beyond the locations of the base stations is required.

2.8.3.2 Signal propagation models

A more sophisticated alternative to simple proximity-based systems is to attempt to infer the distance from a base station from its received signal strength; if several base stations at known locations are visible, multilateration can be used to calculate a co-ordinate position. The main problem with this approach is that signal propagation in busy building environments is very different from what models predict for free space; furniture, doors and people all cause changes in received signal strength that are very difficult to anticipate.

2.8.3.3 Fingerprinting

Fingerprinting is one of the most prominent and successful techniques for building location systems. Since predicting the signal strength from a base station at any given point is very difficult, it relies on building up in advance a radio map of the entire tracking area consisting of measurements of the signal strengths from a number of fixed base stations at each location; to position a mobile target, its measurements of these signal strengths are compared in some way to those at each location on the map and the target is assumed to be at the position whose recorded measurements most closely match its own. This can be very accurate, but has a number of disadvantages, the most significant of which is the ongoing effort required to build and keep up-to-date the radio map.

Even the proponents of fingerprint-based systems acknowledge that this calibration process is likely to hinder adoption. Castro et al. described map-building as “tedious” [26]; Matic et al. rated it “laborious and time-consuming” [142] while Schwaighofer et al. stated “taking calibration measurements is a very costly process, in particular if larger areas need to be covered” [184]. In one of the few large-scale deployments of an RF-based location system, 28 man-hours were required to construct a radio map covering a 12,000 m^2 building [77]. Several techniques have been suggested to mitigate this problem, though each introduces its own issues.

Ocana et al. used a robot capable of autonomously collecting WiFi signal strength measurements in different locations [153]. In addition, they proposed a number of strategies to reduce the calibration effort by optimising the number of collected training samples, thus decreasing the time spent on calibration. This is useful, but a robot is unlikely to be practical in most settings.

Woodman and Harle used an inertial pedestrian tracking system (see Section 2.7.4) to speed up the process of building a radio map for a building by doing away with the need to annotate positions where measurements were taken by hand [208]. Instead, a user walked around the building taking signal strength measurements continuously while his position was tracked. Using this method they constructed a map for a large (8,725 m^2) three-storey building in 2 hours and 28 minutes, travelling a total distance of 8.7 km. This is a significant improvement, but requires an expensive and unusual tracking system and is still an arduous process; data collection had to take place on a Sunday to avoid disturbing building occupants.

Alternatively, maps can be built up by users themselves. Matic et al. developed a spontaneous recalibration mechanism which works by having the mobile device capture a new fingerprint when it is in a known location, such as a docking station or charger, that can be detected by other means, such as the presence an external power supply; the change is then applied to other points based on a radio propagation model [142]. Their evaluation used five defined reference points, and improved the median error of a FM-based system from 1.45 m to 1.2 m after one month's degradation. In practice, this method is unlikely to provide sufficient accurate data points to make a significant difference, and even the example of charging is prone to introduce more error if the phone is ever plugged in in a different location from that configured.

The most practical suggestion is that radio maps could be crowd-sourced directly by asking end users to take measurements in locations that are not well mapped. This is similar to the technique adopted for outdoor WiFi-based localisation by Google for its Android phones, which, every time they fix their position using GPS, also record a WiFi fingerprint and submit it automatically to the central database. *Redpin* is one such system, consisting of software deployed on mobile phones that could capture fingerprints using GSM, Bluetooth and WiFi, though it was not deployed or evaluated on a large scale [19]. Lee et al. used a similar system to evaluate the feasibility of crowd-sourcing a radio map, concluding that it begins to offer reasonable accuracy when the number of fingerprints in the database is larger than 5 for a typical office of 30 m^2 but the recognition accuracy decreases beneath 70% when about 7% of fingerprints are incorrect.

Barry et al. conducted one of the most thorough evaluations of such a system across the five buildings of a college campus over the course of a year, involving more than 200 users [14]. They found that uptake was good, motivated in part by a colleague finding application, with 95% of users contributing and over 1,000,000 location updates in total. They assessed that the system could localise to within 10 m in 94% of cases, although this figure was based only on 57 user estimates of error. These results are very encouraging, but the paper does highlight some significant remaining problems. Changes in the MAC addresses and locations of access points invalidate collected fingerprints and their results may not be indicative of buildings in general since all the users were engineering students or staff who might be assumed to be highly technically literate and all were using institution-issued laptops, mitigating problems with diverse chipsets reporting different RSSI values. Most importantly, initial training is still required to produce a minimally-usable system to which users can be encouraged to add data. There is a chicken-and-egg problem: without good manual training data, users will not use (and therefore train) the system—but without users using it the system will not improve its accuracy. Google addresses this issue in its outdoor system described earlier by recording WiFi fingerprints continuously from the cars used to gather data for Street View. The authors estimated this manual training required 1–3 minutes per location. One improvement suggested is to use shared calendars as a source of training data to help solve the cold start problem [15]. In the same environment they found it yielded similar accuracy to the pure crowd-sourcing approach but in a much shorter time, though it can only be applied in offices where shared calendars with meeting location are the norm.

2.8.4 GSM

Although Place Lab [129] and other systems, including commercial offerings, used GSM towers as beacons to identify a general locality, they were working within the constraints of mobile phones which only reported the cell tower to which they were connected. More recent phones also expose programatically the full list of visible towers and their signals strengths, which has opened up the possibility of using fingerprinting to make more accurate position estimates.

The first such system to promise accurate indoor localisation was introduced by Otasson et al. [155, 195]. Their key idea to make GSM suitable for localisation indoors is the use of wide signal-strength fingerprints including not just the strongest 6 visible GSM towers but all of them, even if they are too weak to be used for efficient communication. They found this could result in up to an additional 29 channels: this higher dimensionality increases localisation accuracy. They collected fingerprints at points located 1–1.5 m apart to achieve a median two-dimensional accuracy ranging from 2.48 m to 5.44 m in large multi-floor buildings; correct floor classifications ranged from 89% to 97%.

Although using a network with almost complete international coverage and a technology that is always enabled on mobile phones is an attractive proposition and some of the systems have attained appropriate accuracy, all of the standard problems with fingerprinting described in Section 2.8.3.3 apply to these systems, making it unlikely that GSM will be a suitable approach to obtain the location information needed by a personal energy meter.

2.8.5 WiFi

RADAR was one of the earliest systems to use fingerprinting based on WiFi signal strength for indoor localisation [10]. It calculated the Euclidean distance between the signal strength vector observed by the target and each entry recorded in the pre-constructed radio map, and produced a best guess of the target's position as the position in the map with the smallest difference. Although this algorithm is relatively simple, it resulted in an error in the original experiments of less than 9 m 95% of the time. Place Lab implemented a very similar system and obtained an accuracy ranging between 3–30 m depending on access point configurations [70, 129]. *Horus* used a different approach, computing a probability distribution of the target's position with a claimed error of less than 1.4 m 95% of the time [210].

Although some of the accuracy measurements claimed by these fingerprinting-based systems are impressive, there are serious obstacles to achieving them in real deployments. WiFi is a power-hungry protocol and using it for continuous tracking will quickly deplete the battery of any mobile device. All of the papers mentioned tested the systems in buildings with far denser deployments of wireless access points than is required just for communications, since it is only necessary to talk to one access point to join a network but a vector of signal strengths from multiple access points is needed to perform matching. Although WiFi networks are almost ubiquitous, most office buildings would not be able to achieve similar levels of accuracy, and the vast majority of homes with WiFi have only a single access point. These are the key locations in which personal energy metering is most important. Furthermore, all the problems described in Section 2.8.3.3 apply; variability of signal strength is a particular concern with WiFi since its propagation is dependent on dynamic factors such as the locations of furniture, doors and even people. A single person between the target and an access point can attenuate the signal by up to 9 dBm [109]. Since the range and the measurement of RSS depends on the wireless card, it is also important to use the same card for collecting the location fingerprints and determining the location; this makes the technique less applicable to devices already carried by users [108].

2.8.6 FM

FM broadcast stations seem a good candidate as the basis for location estimation since they are pervasive across the globe, operate at high power using a very stable reference frequency and, unlike GPS, can be received indoors and in the presence of obstructions. The number of stations visible almost anywhere is a significant advantage; while it may be reasonable to assume that many WiFi access points are visible in an office environment, it is less likely to hold in a rural home. WiFi also shares the 2.4 GHz frequency range with many other electronic devices such as cordless phones and microwave ovens that make it more prone to interference, and its use is prohibited in certain sensitive environments such as hospitals. Furthermore, FM receivers can consume very little power (around 15 mW for one example,[28] compared to 300 mW for WiFi);[29] there are even several solar- or human-powered radios.[30]

Giordana et al. first suggested the use of FM radio signals for localisation beyond simply depending on their limited range so that only devices within range of a particular tower will get information relevant to its coverage area [67]. Krumm et al. introduced the first working prototype based on this idea; the *RightSPOT* system used a vector of radio signal strengths taken from existing public FM stations on different frequencies to allow a wristwatch with a built-in FM receiver to identify its approximate location [124]. However, although they suggested it may be possible to use a signal strength simulator to anticipate the characteristics of different locations without visiting them all, it is still necessary to build a radio map which is a time-consuming and tedious process; furthermore, although FM signals do propagate indoors, the system could only offer suburb-level resolution.

In a later extension, Youssef et al. improved the RightSPOT algorithm to classify a device into a uniform grid of locations rather than identifying one of a few, spatially separated suburbs and validated the simulated signal strength maps suggested previously to eliminate the need for manual training [209]. This allowed the user to be located with median 8 km accuracy.

Fang et al. described a similar system, but performed their experiments in a 1 km^2 area of the campus of the National Taiwan University in Taipei, training the system with 20 reference locations separated by 85–135 m [50]. They compared positioning based on FM

[28] `http://pdf1.alldatasheet.com/datasheet-pdf/view/19431/PHILIPS/TDA7088.html`
[29] `http://pdf1.alldatasheet.com/datasheet-pdf/view/175096/BOARDCOM/BCM4326.html`
[30] `http://www.freeplayenergy.com/`

signal level with GSM and reported that FM measurements provide a lower temporal variation but weaker spatial separation; with the same number of channels, GSM positioning is more accurate. However, FM stations are more widespread, and by using more radio channels and a calibrated spectrum analyser as the receiver, they achieve much better resolution than RightSPOT. With 12 channels, roughly 80% of readings were within 50 m of the true position. For comparison, the Federal Communications Commission requirement for Enhanced 911 to facilitate emergency services is that 67% of position estimates should be within 50 m. However, once again an accurate radio map is required, and the hardware used is significantly more sophisticated than a standard FM receiver.

Papliatseyeu et al. conducted the first experimental study of FM performance for indoor location by employing a dedicated set of short-range FM transmitters built into commercial MP3 players as wireless beacons and a programmable radio on a Nokia N800 Internet Tablet and HMC Artemis smartphone [158]. Their *FINDR* (FM INDooR) system used a radio map and k-nearest neighbour classification of RSSI vectors to yield a median accuracy of 1.0 m, with 95% of readings having a position error of less than 5 m. A comparison with WiFi positioning in the same environment showed a similar accuracy. The authors also shared valuable evaluations of the suitability of various relative position-dependent features and of RSSI stability over time. However, three transmitters had to be installed to cover a single 12×6 m room; each had to be tuned manually to frequencies with little interference from commercial radio stations, and each had a 1.8 m audio cable connected to act as an antenna. Furthermore, a radio map was built based on a 1×1 m grid of cells within a single room; clearly, such a fine-grained data collection would be impractical on a wider scale, and to avoid performance degradation it is necessary to perform periodic recalibration of the system to cope with changing conditions in the environment.

In a subsequent paper, the authors improved the accuracy of their system, introduced a combined, hybrid WiFi and FM approach and introduced a spontaneous recalibration mechanism in an effort to avoid the need to repeat the training regularly [142]. By using a training grid of 0.5 m instead of 1 m (increasing the number of training points from 40 to 140) they reduced the median error by 30 cm. The combined system, fusing FM and WiFi fingerprints, further improves the positioning accuracy by up to 22% to 0.85 m at the 95$^{\text{th}}$ percentile. This approach may capture some of the benefits of both technologies, allowing FM to be employed to provide positioning in areas not well covered by WiFi or to manage battery life. In general, however, FM radio is not well suited to building a location system suitable for personal energy metering; its accuracy is insufficient without extensive and time-consuming radio maps, and almost no devices already carried have appropriate receivers built in.

2.8.7 DECT

Digital Enhanced Cordless Telecommunications (DECT) is a European Telecommunications Standards Institute radio standard for short-range cordless communications; it has an exclusive frequency band with no interference from other technologies that can be used in more than 100 countries with a maximum range of up to 500 m. Its lower frequency allows it to propagate through walls better, and it is permitted 2.5 times the transmit power of WiFi. Being a relatively old and simple standard, DECT chips are comparatively simple and cheap to manufacture; prices for DECT handsets start at around £10.

While not everyone owns a wireless access point, nearly everyone has a cordless phone at home, including those who are not technologically adept. This means the node density of DECT base stations makes it a very good candidate for localisation, particularly in homes and rural areas where it is unlikely that several WiFi base stations will be visible. Kranz et al. reported seeing 5 to 10 times more DECT base stations than WiFi access points across urban, suburban and rural locations in Germany [123]. However, while DECT was standardised in Europe in 1995 it was not standardised in the US until 2005, meaning it is far less widespread there; indeed, their measurements in more than 20 different locations in California did not give a single sighting. Furthermore, very few commercial cards that support it are still available; many have been discontinued in the face of the rise of WiFi and VOIP.

Schwaighofer et al. presented an probabilistic algorithm for localisation in generalised cellular networks and evaluate it using DECT [184]. The key idea was to use Gaussian process models for the signal strength received from each base station, and to obtain position estimates via maximum likelihood. Again, a radio map was required; in a 250×180 m assembly hall, measurements were made at 650 points. The total number of base stations installed was not specified, but localisation was typically based on around 15; this is rather higher than would be expected in most homes or offices. Using the full set of calibration measurements, their algorithm performed slightly worse than k-nearest neighbours, with an average error of 7.5 m; its strength is that it performs well with less calibration, achieving an average error of around 17 m with only 12 calibration points 75 m apart (kNN achieves 29 m at best). This is a significant advance, and many subsequent systems adopt the idea of Gaussian processes, but the use of DECT was incidental to the evaluation of the algorithm and it does not really offer any information about its merits and demerits relative to other radio technologies.

Kranz et al. built the first real-world system based on DECT and performed a comparative study of DECT and WiFi signals for indoor localisation [123]. It used fingerprinting and a weighted kNN method in a very similar way to many early WiFi-based systems; additional accuracy could be obtained by the use of more sophisticated probabilistic algorithms. They evaluated it in urban, suburban and rural settings. In an apartment in a residential urban area both DECT and WLAN fingerprinting performed very poorly, with results comparable to selecting a location at random. However, in a suburban office, using 24 reference points surveyed with sub-centimetre accuracy but no artificial DECT stations, more than 55% of the location estimates showed less than 5 m error and more than 90% less than 10 m error. This was slightly better than the equivalent WiFi figures. Similarly, in the rural setting, DECT outperformed WiFi, with 80% of all estimates showing less than 5 m error.

As a technology, DECT seems well-suited to positioning, but there are significant obstacles to a global system. Mobile phones and wireless networks mean its days are likely to be numbered, though it may have a part to play in a hybrid system. More importantly, very few mobile computing devices that people carry with them each day feature DECT; it is designed for a cordless phone that remains within a single building. While it might therefore work well for context-aware applications within a home or office environment, it is less suitable for a personal energy meter that requires the device to be able to position itself indoors in a number of different locations; a separate, dedicated device would have to be carried.

2.8.8 Bluetooth

Madhavapeddy and Tse conducted an extensive survey of Bluetooth signal strength in the William Gates Building [137]. The Bat system was used as a 'location oracle' to gather a large number of samples of Bluetooth signal strengths from transceivers placed around a building. They concluded that Bluetooth was ill-suited for the purpose of accurate, low-latency location sensing due to:

1. common chipsets only exposing a running average of signal strengths and updating it infrequently;
2. the high variance in signal strengths for longer distances; and
3. the inability for consumer mobile phones to maintain multiple Bluetooth connections simultaneously for triangulation purposes.

Their third point has been addressed by most modern mobile phones, but it is the second one that poses the most significant problem: most other authors of papers concerning Bluetooth location agree that the coarse RSSI metric is not a good indicator of distance [13].

Most Bluetooth tracking systems are proximity-based — that is, if a user can be contacted by a base station (or vice versa) then the user is coarsely located to the base station's position. Bluetooth was designed as a short-range communications system with range of comparable size to a room (class 2 Bluetooth devices have a nominal range of 10 m in free space, less indoors), so proximity-based location is simple to implement and sufficiently accurate for many purposes, including energy metering.

Anastasi et al. conducted some of the first experiments to demonstrate the potential of Bluetooth as the basis of an indoor location system [8]. Their *Bluetooth Indoor Positioning System* (BIPS) is proximity-based and outside-in; users must register their user ID and Bluetooth device address with a centralised server that then orchestrates a network of fixed base stations to track the movements of mobile devices using the Bluetooth inquiry mode [23].

Huang implemented a proximity-based, inside-out system, where Bluetooth devices are placed in key locations throughout a building and respond to enquiries from a mobile user's phone or PDA [100]. The work has a refreshing emphasis on practicality: existing PCs were used as beacons, with USB Bluetooth adaptors installed in computers approximately every 10 m on six different floors. An advantage of the inside-out approach is that no software beyond the device drivers is required on the fixed PCs; Huang asserted:

> On average, configuring a machine to host one of our beacons took less than three minutes. The most time consuming part of the deployment was actually tracking down the system administrators for the machines we wanted to use, and obtaining their permission.

Client software was written for both Linux PCs and Symbian mobile phones. The main obstacle encountered was the time taken to perform a Bluetooth enquiry (up to 10.24 s); to help get around this, Huang installed two separate Bluetooth adaptors in each beacon

PC instead of one. Experiments showed that they responded to enquiries independently from each other, and the locator only needs to wait for a response from either one. Unfortunately, this doubles the cost of deployment.

In order to remove the dependence on host PCs outside his control, Huang also evaluated the possibility of deploying dedicated low-cost beacons constructed by connecting a USB Bluetooth adaptor directly to a powered USB hub. A PC is then only required to be connected to the hub to initialise the adaptor, after which it can be removed. Unfortunately, this method requires significant manual effort to re-initialise every adaptor in the event of even a momentary power loss. While this may not be an insurmountable problem in a Computer Science laboratory, it is likely to be impossible in a general deployment.

With any proximity-based system it is clearly desirable to use beacons with a range approximately equal to the required spatial resolution. Cheung et al. built on Huang's work and proposed a new type of dedicated beacon, constructed from a Bluetooth headset placed inside a cardboard carton and wrapped in foil tape to reduce the range to approximately 2–3 m [27]. This does not suffer from the initialisation problems of Huang's solution. The short range improves resolution, but would require a very large number of beacons to offer complete coverage. Removing the headset battery and leaving it connected to mains power meant it remained permanently in standby mode, where it was not discoverable but would respond to page attempts to its own address. They therefore built a new test client that instead of performing Bluetooth scans continuously attempted to connect to the addresses of known beacons; each attempt took between 1.5 and 6 s. Clearly, a naïve algorithm like this will not scale to a large building-wide deployment with hundreds of beacons.

Naya et al. proposed a Bluetooth-based proximity detection method in a nursing context [149]. They too relied on the Bluetooth discovery mode, but by an exhaustive search through the inquiry parameter space they reduced the mean turnaround time of inquiry responses to less than the expected time of the Bluetooth default settings at the cost of some reliability. The method only appears to have been tested with a single master and slave, and there is no record of subsequent development and deployment of a working system.

Hallberg et al. implemented an inside-out, proximity-based system with two separate mechanisms for resolving a visible Bluetooth device to an absolute position: either the mobile client could look up its address in a central database, or it could connect to a special service running on the fixed target and ask it directly [79]. The latter approach has the advantage of not requiring any additional network connection or any co-ordination between target devices, which can be completely independent provided they know their own locations, but clearly only works if the targets are PCs running custom software. The client performs a Bluetooth inquiry, then establishes a connection to each discovered device to ask its position; the position estimate is a simple geometric combination of the positions of the devices in range. The time taken to arrive at a position estimate therefore scales linearly with the number of devices.

Bluetooth tracking has also been demonstrated using the received signal strength measurement techniques first developed for WiFi location systems [76, 157], based either on radio propagation models or 'fingerprinting'. For example, *Bluetooth Local Positioning Application* (BLPA) was an attempt to use received power levels to infer distance from base stations using a radio propagation model learned from a calibration process [122]. There are significant problems with this approach, not least of which is that the model will be subject to continuous change and RSSI readings are dependent on the particular chipset. Although the authors claim sub-room-level accuracy, the system was only tested in a single room with all WiFi devices disabled to avoid interference. Genco et al. adopted a fingerprinting approach in a castle in Sicily and noted the difficulty in predicting signal propagation [66]. They proposed methods for iteratively determining the optimal positions of base stations and achieved an apparent accuracy of around 0.5 m using 10 base stations, but this relied on techniques specific to the castle in question that they admitted would not be generally applicable.

Bargh and de Groote took a different approach and demonstrated a fingerprint-based solution that relies only on the response rate of Bluetooth inquiries [13]. Target devices must be discoverable; fingerprints are based on the percentage of responses to inquiries sent by dongles that are located at a specific distance from each other. This has the advantage of working even if different devices are used for the calibration and localisation phases; it produced an accuracy of around 98% correct room estimation with full overlapping Bluetooth coverage and around 75% with only partial coverage. However, the experiment involved leaving the target device in each position for an hour; this will always produce better results than attempting to track a moving target; furthermore, the authors acknowledged the method for obtaining the fingerprints is infeasible when the number of locations is high.

Although signal-strength-based systems generally claim greater accuracy, that comes at a significant cost; all of the problems with these methods described in Section 2.8.3.3 apply in addition to problems caused by using the inquiry scan described in this section. Since the resolution offered by proximity-based systems is sufficient for most personal energy metering applications and the variability of Bluetooth signal strengths over time and with different chipsets is well known [137], systems based on signal strength are discounted here.

Bluetooth's short range makes it the most promising technique for indoor location that does not require fingerprinting or other extensive calibration processes; the systems described above demonstrate the potential of this technology. However, their common requirement for devices to be permanently 'discoverable' is a major concern. This problem, and a possible solution, is discussed in Section 5.1.1.

2.9 Syndication of sensor data

The personal energy meter will need to aggregate data from a number of different sensor systems to present a unified view to its end users.

Pachube[31] is a data brokerage platform for Internet-connected sensors, which can provide data in either CSV or an XML feed format.[32] The site acts as a central point, storing all of the data collected and serving it to consumer applications; this model has proved successful for sites such as YouTube, from which Pachube takes inspiration, but has the same disadvantages as all centralised systems described earlier.

Guinard and Trifa noted the popularity of mashups on the web and discussed how the REST principles [52] can be applied to embedded devices, either by implementing a web server directly where resources permit or via a gateway [72]. As part of the *Web of*

[31] `http://www.pachube.com/`
[32] `http://www.eeml.org/`

Things project they showed how an eco-system of these devices can facilitate the creation of 'mashups' consuming real-world data, but noted that REST can complicate the creation of certain more complex composite services.

sMAP is a design for RESTful web services to allow sensors and other producers of physical information to publish their data directly [40, 41]. It adds a layer of sophistication to simple syndication formats, including support for actuators and basic transactional semantics. The evaluation deployment in a building at the University of California, Berkeley, provides about 2,000 measurement channels monitoring electricity consumption, environmental quality data, HVAC parameters and weather data. The system is intended for a different level of operation from the personal energy meter, allowing separate sensor systems within a building to communicate with each other; as such, it promises to be a good basis on which to build individual components that report to a personal energy meter.

To demonstrate the potential of sMAP the team created what they call a *Human-Building-Computer Interaction* (HBCI) system, consisting of an Android mobile application, a number of RESTful services implemented using sMAP and 2D barcodes placed on physical objects [99]. By scanning the tags with a mobile phone's camera users can interact with physical objects, switching appliances on and off and viewing graphs of their current and historic energy consumption. Although its triggers demonstrate the support sMAP offers for actuators they are unlikely to prove convenient for most people, but the use of barcodes to identify objects or spaces of interest is a straightforward and easily understood mode of interaction that is also adopted by the example aggregator described in Chapter 6. HBCI also includes preliminary support for apportionment developed from the principles and ideas set out in Section 3.5. Overall, HBCI is a good demonstration of several ideas set out in this dissertation.

2.10 Summary

Personal energy metering stands at the intersection of a number of disparate strands of research, from wireless sensor networks to persuasive technologies and location systems, and it is necessary to consider established work in these fields and how it could fit into the overall architecture. This chapter has shown that feedback is a valuable tool for reducing energy consumption; based on results from previous projects, it should be understandable at a glance, show progress towards goals and allow for easy comparisons. Many systems already exist for metering consumption in each category, and these could all feed into a personal energy meter, but few attempts have been made to infer consumption using minimal sensing. Similarly, although many indoor location systems have been suggested and location will provide an invaluable form of context for energy metering, all require investment in either equipment or calibration that makes it unlikely they will be deployed solely the purpose. The remainder of this dissertation therefore focusses on methods to derive the necessary data in an incremental fashion.

Chapter 3

Methodology and apportionment

> The... most mysterious piece of nonabsoluteness of all lies in the relationship between the number of items on the check, the cost of each item, the number of people at the table and what they are each prepared to pay for.
>
> (The Hitchhiker's Guide to the Galaxy)

Contents

Overview

This chapter demonstrates the importance of apportioning energy costs to individuals and the potential of sentient computing technologies to contribute crucial information. It enumerates the concepts of apportionment, including a taxonomy of resources and the data requirements to handle them, evaluates a number of different apportionment policies and derives the principles that are generally applicable: *completeness*, *accountability* and *social efficiency*. It describes a method for estimating building occupancy and the technology required to estimate each form of energy consumption. It also presents the methodology used for the remainder of this dissertation.

3.1 Methodology

An *energy stack* adds up our energy consumption in its various forms and represents the answer by a stack of blocks whose height is proportional to the energy used. This

Some of the contributions presented in this chapter have also appeared in separate publications [86, 90]. Figures 3.2 and 3.4 are reproduced courtesy of Andrew Rice.

Figure 3.1: Estimated energy consumption of a "typical moderately-affluent person"

allows for immediate visual comparisons. Figure 3.1 shows MacKay's estimate of the consumption of a 'typical moderately-affluent person', which comes out to a total of 195 kWh per day [136]. The personal energy meter should be able to refine this estimate incrementally to produce a personalised and dynamic energy breakdown for an individual. The estimates provided by MacKay's stack also provide a systematic basis for identifying which areas of consumption are most significant and therefore merit most effort.

Although future buildings might have sophisticated management systems, future vehicles might make available all their sensor data and that everyone might carry smart devices which provide a wealth of data from which to derive accurate personal energy accounts, it is necessary to consider the current situation in which sensor network deployments are sparse and unreliable. A successful personal energy meter must make worthwhile estimates based on very little hard information which can be progressively refined as more data becomes available. Clearly, if users are to be billed for their energy consumption then dependable and definitive data is a necessity, but if the objective is to provide a feedback tool to help users who are already committed to reducing their footprint then absolute numbers are less important than scales and trends. It is unlikely that anyone will deploy expensive dedicated sensor systems just to drive energy metering, so the output of any such scheme must adjust as the input improves from zero to total knowledge; between the two extremes will lie a sweet spot that offers the best compromise between cost and utility. The theme of *incremental sensing* runs throughout this dissertation: as discussed in Section 1.4, it should be possible to start with minimal data and increase the quality of the result by adding additional sensing.

With no sensor data or other inputs, the best that can be offered is an estimate based on aggregated knowledge about the habits and trends of the entire population—such as in MacKay's stack (Figure 3.1). Given answers to questions about lifestyle and habits, it is possible to personalise the estimates in the stack to an individual; this is the approach adopted by a number of 'carbon footprinting' services on offer today.[1] This remains a static allocation because changes in lifestyle require updating old answers, but it does provide a foundation on which to build. Individual segments can be incrementally replaced with live data as it becomes available, while others remain an estimate.

The remainder of this chapter discusses in detail how these estimates can be obtained, working from top to bottom through MacKay's stack.

[1] such as `http://wattzon.com/`

3.2 Transport

3.2.1 Jet flights

Air travel is relatively easy to record by monitoring email communication of airline bookings. This technique often forms the basis of the carbon footprinting calculations made by online tools, as described in Section 2.4.2. More sophisticated data mining techniques could factor in airline data on seat occupancy and fuel burn to improve estimates of each passenger's share of the energy for a flight.

Unlike the embodied energy in goods which should be amortised over the expected lifetime of the item, flights have a definite duration over which the energy should be allocated.

3.2.2 Car

The locations of bus and railway stations can be combined with a GPS trace of journey start and end points to estimate mode of transport and hence energy consumption. However, distinguishing between travel by foot, bicycle or car is much more difficult in a congested urban environment. Additional data from inertial sensors (now common in many phone handsets) might help with the classification problem; see Section 2.4.2 for a survey of possible solutions. Where several people share a car, the energy cost for that journey should be split between them. This requires contextual information on location or identity to determine who is traveling in a given vehicle (see Chapter 5).

3.3 Public services

It seems reasonable to divide the energy cost of essential public services amongst the population, in the same way that these services are funded through general taxation. This is a specific case of the general principle that energy allocated to someone for work done on behalf of others should be *delegable*, in the same way that the lecturer printing thousands of pages of course notes should not be held personally responsible for the energy costs, but should be able to pass on a small fraction to each undergraduate he teaches.

3.4 Buildings

In 2004, buildings accounted for 37% of total energy consumption in the EU.[2] This figure includes heating and cooling, 'gadgets' and other electric devices and lights. The planned national roll-out of smart meters in every home will commodotise the collection of live data on household consumption. Devices aimed at the consumer have recently reached the market which monitor and record total electricity consumption; these are discussed further in Section 2.3.1.1. Andy Stanford-Clark has for many years has published real time graphs of his house's electricity and gas consumption on the Internet.[3] He explains that although he had to built his own hardware to achieve this, the emergence of this batch

[2] International Energy Agency. Key world energy statistics. 2006.

[3] `http://stanford-clark.com/power/`

of packaged devices has led to hundreds of his colleagues implementing similar systems in their homes. Nevertheless, these devices are not yet widespread; most individuals still have old-fashioned meters, and a personal energy meter will have to cope with anything from sporadic manual readings to live up-to-the-minute streams.

MacKay breaks down his estimate for building energy use into home, work and public sector. In order to make use of the collected data for an office or other shared building it is necessary to decide what proportion of its energy demand is attributable to the individual in question. This process is referred to as *apportionment*. As discussed in Chapter 2, there is little discussion in the literature of methods for disaggregating consumption by person, so additional studies were conducted. Section 3.5 evaluates the differences between a number of possible policies, while the examples in Section 3.6.2 demonstrate the importance of apportionment.

3.5 Apportionment

Apportionment is defined by this dissertation as the process of dividing the total consumption of a building, organisation or other entity and allocating it to individuals in proportion to their use. A number of different possible strategies can be considered. The most obvious of these is simply to apportion a static fraction of the building's power consumption to all those who work there—and such a strategy might well be suitable for a private residence. However, for larger buildings this is a poor solution. Lutzenhiser argues that overlooking the variability of human social behaviour "significantly amplifies and dampens the effects of technology-based efficiency improvements" [135], while 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 [68]. 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. Here lies the tragedy of the commons [81]; if the incremental cost of an action is always less than the expected benefit because of the apportionment scheme in use it will tend to lead to undesirable behaviour.

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 to overcome in order to achieve an appropriate level of sensor coverage to provide this information.

3.5.1 Apportionment policies

Dynamically varying the proportion of the building's consumption assigned to each individual allows a policy to capture the variation in energy due to their activity. There are numerous possible policies to determine how it should be carried out and different ones

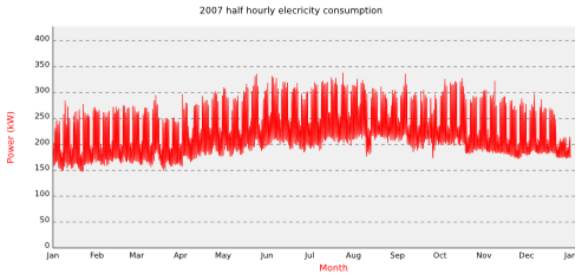

(a) Electricity consumption plotted against time

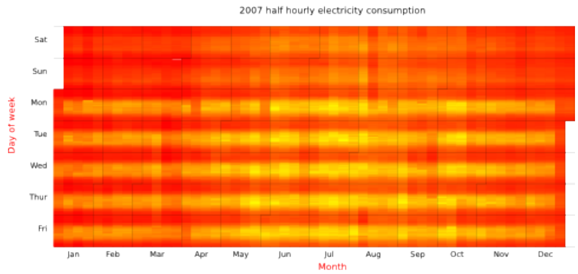

(b) Heatmap showing weekly and seasonal patterns in electricity consumption. Brighter colours indicate higher consumption.

Figure 3.2: Half-hourly electricity consumption for the William Gates Building in 2007

will suit different buildings and organisations. Nevertheless, there are certain desirable properties that all apportionment policies should exhibit:

Completeness: the sum of the energy apportioned to all individuals should be equal to the total energy to be apportioned

Accountability: actions by an individual should have a maximal effect on their own allocation and a minimal effect on others

Social efficiency: actions taken by an individual to reduce his allocation should also reduce the total energy used

To evaluate the differences between a number of possible policies a separate case study was carried out based on energy and usage data for the William Gates Building[4] collected over the course of a year (Figure 3.2). Initial estimates were iteratively refined through the addition of further sensing. The remainder of this section describes the details and results of this study.

The result of apportionment is necessarily specific to a particular individual and three hypothetical working patterns based on representative individuals for the building in question were considered: 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 Table 3.1. In this section, a variety of apportionment policies are described and evaluated with respect to these three individuals for a typical week in November 2007. The total apportioned energy for each policy is given in Table 3.2; these are explained in more detail in the following subsections.

3.5.2 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 the building is around 250 and total energy consumed in 2008 was 2,025,778 kWh, meaning a

[4]The William Gates Building was opened in 2001 as the new home of the University of Cambridge Computer Laboratory. For more information on its architecture, see `http://www-building.arct.cam.ac.uk/westc/cl/cl.html`

	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

Table 3.1: Working patterns of example individuals

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

Table 3.2: Total energy (kWh) allocated by the apportionment policies for a week in November 2007

user allocated a $\frac{1}{250}$ share would be responsible for 8,103 kWh. For comparison, the total energy consumption for one participant's house for the same year was around 2,200 kWh.

The electricity meter of the office building, in common with those of many large buildings, logged half-hourly measurements of the total energy consumed. The resulting power apportioned for the example week is shown in Figure 3.3. The line is a scaled version of the overall power consumption (Figure 3.2(a)); all users pay more on weekdays, regardless of whether or not they were present. This policy violates the principle of accountability by making no accommodation for working patterns or individual actions—any and all consumption is shared amongst all building users.

3.5.3 Dynamic apportionment

To improve upon static apportionment it is necessary to take into account the behaviour of individuals. In the first instance it is assumed that all building users behave similarly and so apportionment is based on the current occupancy of the building.

As one might expect, building occupancy varies significantly over time. Low occupancy is expected over weekends and holiday periods but also due to less predictable causes such as

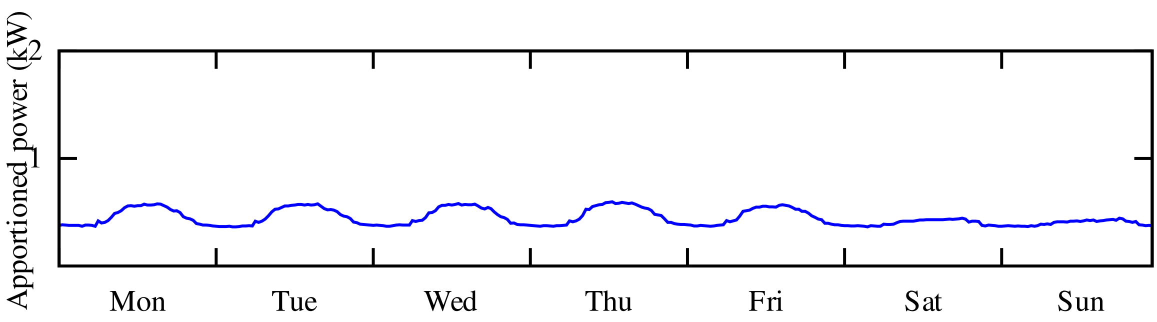

Figure 3.3: Power apportioned to each individual under the 'equal' policy

travel disruption—for example, the UK transport infrastructure copes poorly with snow. For this reason some form of sensor data will be required for occupancy estimation.

A variety of sensor systems could be used to provide this information, including fully fledged location systems, 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 to make use of systems that are in place today before adding more sensing. Chapter 5 provides a full discussion of the importance of context and possible sources of contextual cues and presents a novel low-infrastructure location system designed to provide input to a personal energy meter.

3.5.3.1 Estimating occupancy

Although some parts of the building in this case study 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, it is not possible use these to infer who is in the building at any given time. However, the logs can provide 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 these purposes.

Under the assumption 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 occupancy the following algorithm was used:

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.

Figure 3.4 shows the estimated occupancy trace for 2008. Dark colours correspond to a large number of people in the building and lighter colour correspond to fewer people. It shows quiet days which correspond to UK public holidays and a general ebb and flow corresponding to term and vacation periods. Note there are a number of exceptions such

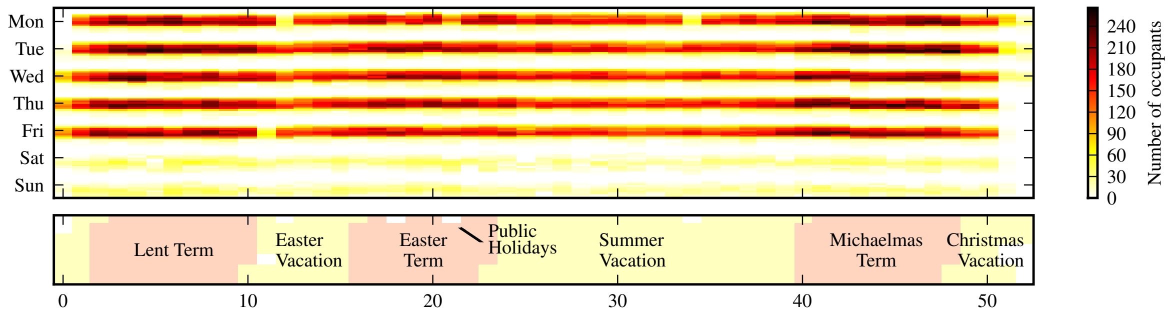

Figure 3.4: Estimated occupancy trace for the William Gates Building for 2008

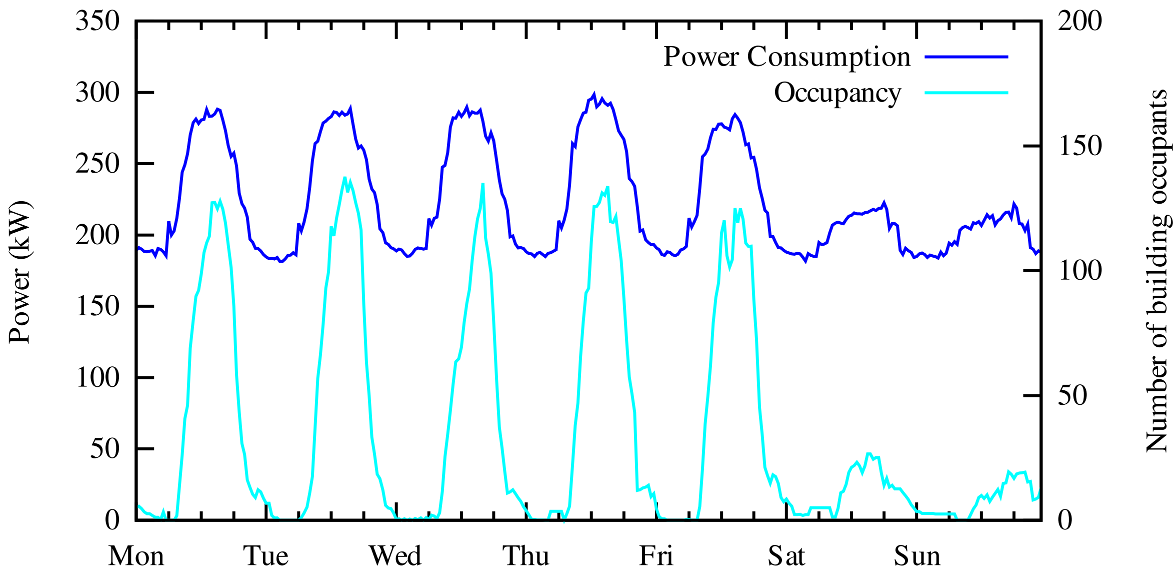

Figure 3.5: Estimated occupancy and power usage for the William Gates Building

as public holidays during term time (which the university does not observe) and the especially quiet period for the two weeks surrounding Christmas.

The estimates this produces for the example week are illustrated in Figure 3.5. Estimated occupancy and power usage are strongly correlated; the occupancy drops off dramatically at weekends, and dips at lunchtime are clearly visible.

3.5.4 Occupants policy

The 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 the example week are shown for several typical working patterns as the dark lines in Figure 3.6. 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, the 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.

3.5.4.1 Base load

To improve on this policy, the base load can be estimated and divided 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 3.6. 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 also better because 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 evenly 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 computers and lights on overnight—this results in a *lower* energy cost to the individual than switching them off, since they are then included in the base load and split between many more people.

3.5.5 Personal load policy

Instead of estimating the base load and assuming the remainder is personal, the problem can be approached 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 the offices with a simple power meter revealed that the devices everyone typically switches on when they arrive, such as lights and monitors, consume between 100 and 200 W, depending on office size and computer configuration. Supporting this observation, the 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 the earlier estimate of the base load. Figure 3.7 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 one might 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 has been considered so far. Instead of simply dividing up the

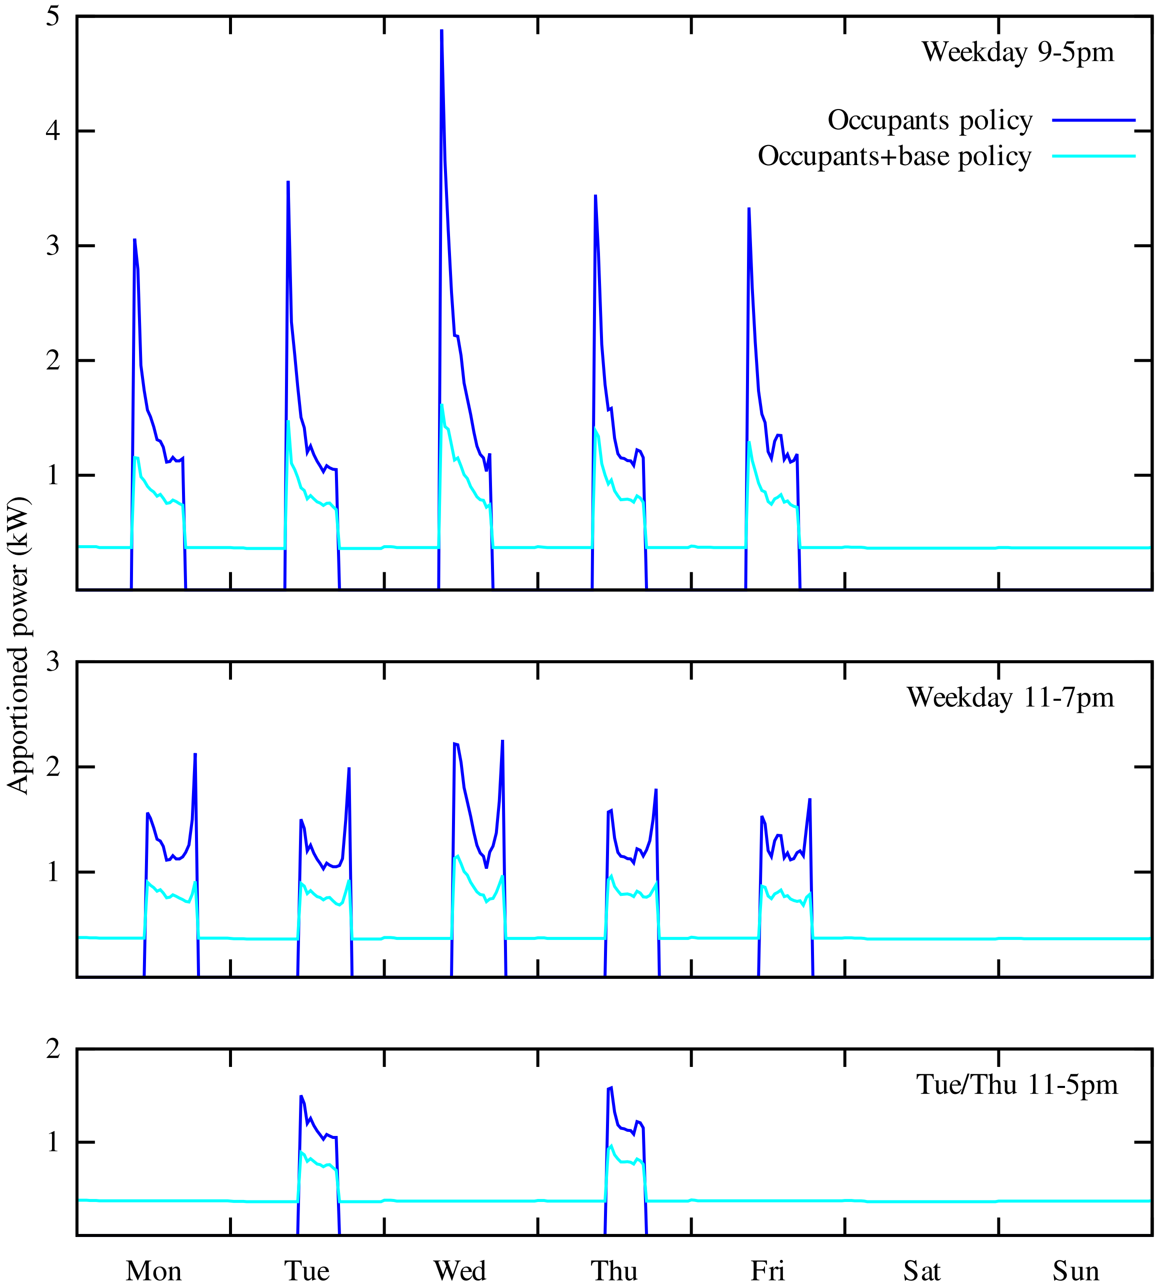

Figure 3.6: Power apportioned under the ‘occupants’ policies to example individuals (1: top, 2: middle, 3: bottom)

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. This problem is addressed in Chapter 4.

Clearly, the different policies make a significant difference to the end results, but it is interesting to note that the ‘occupants+base’ and ‘personal load’ policies produce broadly comparable numbers. This supports the intuition that both are reasonable strategies and the only difference between them is a result of an inaccurate estimate of the base load: with omniscient sensors that could tell exactly which devices were consuming power the two policies would become the same.

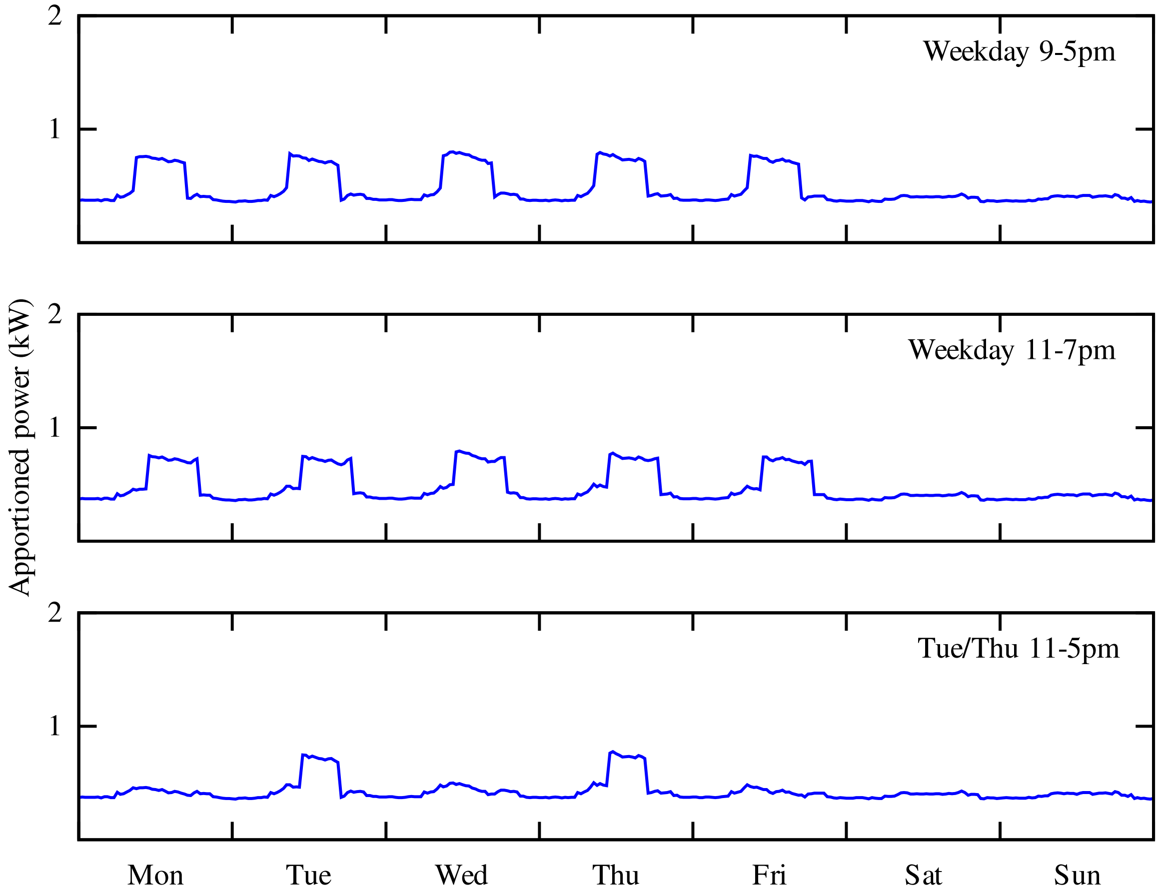

Figure 3.7: Power apportioned under the 'personal load' policy to example individuals (1: top, 2: middle, 3: bottom)

3.6 Gadgets and 'stuff'

The overall energy consumption trace for an appliance (Figure 3.8) consists of a) the embodied energy which is incurred during manufacture; b) energy due to usage; and c) recovered energy from the return or recycling of the device.

A large proportion of the typical energy stack represents the *embodied energy* of the things we buy or use; this includes raw material extraction, manufacture, assembly, installation, disassembly and deconstruction. A further significant contributor is the energy required to transport 'stuff' from the point of manufacture to the point of use; MacKay also separates food, farming and fertilizer from other classes of product, but similar principles apply to all. Much like the principle of depreciation in accountancy, the embodied energy of a product can be amortised over its expected lifetime—though this allocation may change if the product is subsequently shared with others. Knowledge of items purchased could come from an inventory which would also provide profiles detailing their embodied energy, expected lifetimes and other users. Looking ahead, barcode reading applications are now included in many camera phones, while more and more products are shipped with RFID tags, offering opportunities to catalogue purchases automatically.

One model for apportioning usage is apply the direct costs to the current user but to share the indirect energy costs between all possible users of the appliance. This means that an individual's energy bill will reduce as additional users of the appliance are registered and

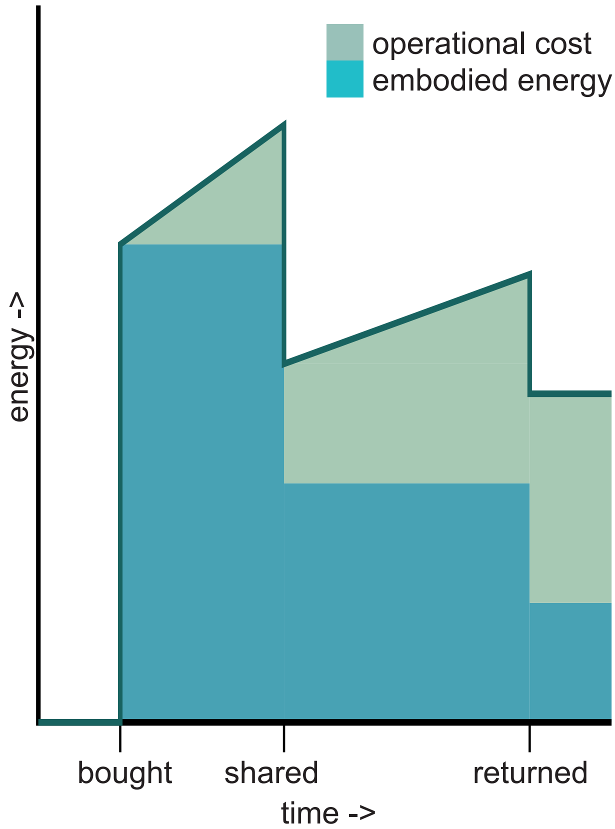

Figure 3.8: Total energy cost of ownership

thus take a share of its energy cost. A 'refund' of embodied energy could be credited if the device is recycled; alternatively, if it is thrown away before the end of its anticipated life, the remaining embodied energy not yet allocated must be accounted for.

For each resource, there is a set of people who have access to it and a subset who are using it at any given time. These sets are likely to be needed by most apportionment policies to allocate shares of the energy costs appropriately. For ease of discussion, it is possible to categorise certain cases of set membership and thereby categorise resources:

Owned resources are those which are only used by a single person who should always bear the responsibility 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. No contextual information is needed to allocate their energy costs; a simple inventory is sufficient, though building and maintaining this inventory is not an insignificant undertaking.

Shared resources are those to which several people have access but which are only used by one person at a time; examples are printers, photocopiers and showers. It may be appropriate to share their embodied energy and baseline costs amongst all potential users but to allocate their instantaneous consumption to the user directly responsible.

Communal resources are those which benefit several people simultaneously: examples are heating and lighting, or public transport. It may not be reasonable to divide their instantaneous consumption amongst those using them at the time: energy spent on heating a building while nobody is in it benefits those who arrive later.

3.6.1 Owned resources

To apportion the energy costs of owned devices requires only an accurate inventory associating devices with owners and, if direct metering is not possible, with appropriate power profiles.

3.6.2 Shared resources

The obvious mechanisms for handling shared and communal resources are either simply to divide their total energy cost amongst all those entitled to use them (static apportionment) 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; additional sensors are required 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, it is necessary to measure the energy consumed in each individual interaction. Device profiles must identify the energy costs of specific events, such as printing a page or making a cup of coffee. This is discussed in more depth in Chapter 4.

3.6.2.1 Printing

Printing provides a good case study of the value of apportioning use of shared resources since print server logs provide second order information on usage. This allows accurate analysis without the need for the sensor systems that would be required to determine usage for most other equipment.

The printer logs for the building were analysed. The logs cover 28 printers for a period of 47 days, during which time 198 users printed a total of 82,349 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 3,452 pages, while the lightest printed just 1, and the top 15 users accounted for over half the total printing between them. The average number of pages printed was 416, and the mean percentage deviation from this average was 86.5%, highlighting the importance of profiling and a dynamic apportionment policy rather than a simple static division of the printer's energy usage.

Energy measurements can be combined with these logs to improve on the 'personal load' policy; see Section 4.4.2 for details.

3.7 Summary

This chapter has ascertained the categories of consumption that are significant and the inputs required to build a personal energy meter using incremental sensing: *metering* and *context*.

The simulation of several policies for each participant has shown that apportionment is important and the correct choice merits careful consideration. Different policies have significant effects on the total energy allocated to individuals, but all should abide by the general principles of *completeness*, *accountability* and *social efficiency*. 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 (Chapter 4) as well as low-infrastructure identity and location sensing systems (Chapter 5).

Chapter 4

Modelling and profiling

Contents

Overview

This chapter shows how computing systems can be used to measure or calculate energy consumption. It demonstrates that device profiles and inventories can be combined with secondary indicators of activity to infer consumption without live metering (Section 4.1), and presents a novel method for modelling building energy consumption based on profiles and crowd-sourced inventory data (Section 4.2). It then explores how these device profiles can be created (Section 4.4) and describes a new technique for decomposing power measurements of programmable devices to profile them in an automated manner (Section 4.5); this is necessary to apportion the energy costs of shared resources.

4.1 Profiling and secondary indicators

The previous chapter demonstrated how a building's energy consumption can be apportioned among its users to present personalised feedback to each individual. While this apportionment is important, it is of limited use to a user to know only his or her total consumption; without a breakdown by function, it is difficult to identify where changes could

Some of the contributions presented in this chapter have also appeared in separate publications [177, 175, 176]. Dan Ryder-Cook created the physics model described in Section 4.2. Brian Jones assisted with the design of the circuit boards shown in Figures B.1 and B.2. Ee Lee Ng recorded the LCD data in Section 4.4.3.

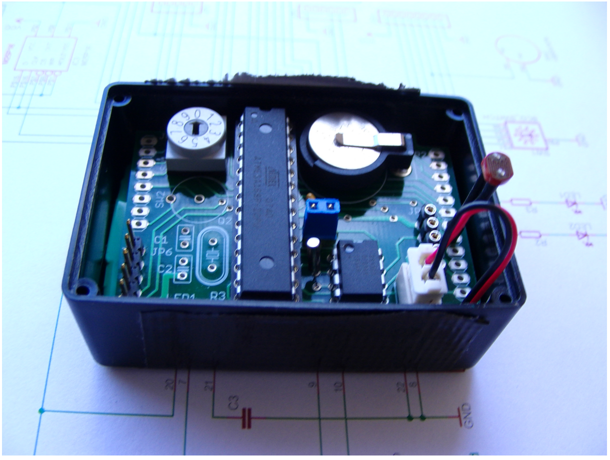

Figure 4.1: Sensor node built to record lighting state

be made. This chapter therefore presents mechanisms for disaggregating consumption when only limited data is available.

Live metering of energy consumption is becoming more prevalent. Taking electricity as the most common example, Section 2.3 describes some of the myriad technologies now available to deploy. It is easy to imagine homes and offices of the future in which every socket is 'smart' and reports the consumption of the device attached in real time over a common network—perhaps even over the power line itself. Plugs might contain tags to identify the device in question, allowing an inventory of a building to be created automatically. Such an infrastructure would be a boon for personal energy metering, but it remains some way off. In the interim, circuit-level submetering is becoming more prevalent, but this provides a spatial, rather than functional, disaggregation. Although the two may be roughy equivalent for some centralised loads, such as server rooms or heating, cooling and air conditioning, this is not generally the case.

For a personal energy meter to be useful it must be able to function without ubiquitous sensing; early adopters should derive benefit from its results even with very little input data available. This chapter therefore investigates ways in which energy consumption can be inferred or estimated rather than measured directly and continuously. One such mechanism is to *profile* a device, enumerating its power states and the energy costs of performing each possible action. Assuming devices have regular, deterministic profiles, a profile built for one device can be applied to others of its type in a building inventory, and secondary indicators, such as cheap or repurposed sensors, logs or even human input can be used to determine which state each device is in or which actions it has performed—thus removing the need for continuous metering (see Section 3.6.2).

This method is well-suited to apportionment, since it makes clear the costs of individuals' actions rather than the cumulative costs of devices, and also aids the evaluation of hypothetical scenarios, answering questions like "what would be the saving if PCs were switched off overnight?"

4.1.1 Low-fidelity estimates of lighting energy demand

Lights are a good example of a resource that often cannot easily be metered; the wiring for ceiling-mounted lights in office buildings is generally inaccessible except to electricians and several rooms or corridors share a single lighting circuit. Fortunately, it is easy to infer their consumption using profiling and secondary indicators: most lights have only two states—'on' and 'off'—and a known, static power usage, while light sensors cost pennies and are readily deployable. Provided it is known which lights are on which circuits, it is not even necessary to instrument each light; knowing that one is on is enough to know that all the others on the same circuit are on as well.

This was demonstrated by *LightWiSe* (LIGHTting evaluation through WIreless SEnsors), a wireless tool which aimed to evaluate lighting control systems in existing office buildings [42]. The experiment used a set of TinyOS-based motes, each with an on-board light sensor used to detect ambient light and luminary state and an additional sensor board containing a long range PIR sensor used to detect occupancy. The authors concluded that significant savings could be made from turning off lights automatically when they are not required.

A manual survey revealed the number and model of lights installed in each room of the William Gates Building; the manufacturers' specifications include their power usage. This data can also be 'crowd sourced' by asking users to mark the lights themselves [178]—see Section 2.5.2 for details.

A small and cheap microcontroller-based sensor node was developed which can be attached to a light with sticky tape and will log to internal Flash memory the times at which the light was switched on and off (Figure 4.1, and Figure B.1 on page 180). A simple light dependent resistor and binary threshold suffice if the sensor is mounted directly on the light, since the increase in light levels when it is switched on dwarfs any variation due to sunlight coming through the windows.

Although the nodes were powered by batteries that will last many months between charges, to ease deployment and maintenance costs, such devices could be powered by harvesting energy from the lights themselves in a similar way to that used by the Locust infrared location system [118]. This approach was not adopted both to speed up design and development and to keep costs low. These sensors were deployed in the offices in the research group; in each room, there was a single switch or motion sensor that controls all of the lights so knowing whether one light is on is sufficient to infer the energy consumption of all of them. The lights in corridors operated on a timer, so no sensing was required; Building Services provided the details of all automatically-controlled lighting. To verify the correct operation of the sensors, a manual log of the times the lights were switched on and off was also kept in one office for comparison.

Most lights are in individual offices and can therefore be classed as owned resources and their power allocated to the owner whenever they are on. Others, such as external and corridor lighting, are a public service in the same way as heating and certain computers; they benefit everyone connected to the building and can be counted as part of the base load. The most interesting case are those in shared offices, where it is necessary to combine knowledge of the lighting state with occupancy data to apportion a share of their power to each office owner who is present while they are on.

The sensors provided an accurate log of when the lights were on, but in keeping with the principle of attempting sensible estimates from limited data which can then be refined as more detailed information becomes available, the possibility of estimating the state based on knowledge of external light levels was investigated. Every office in the building in question has external windows; in offices without, which can be determined from Open-

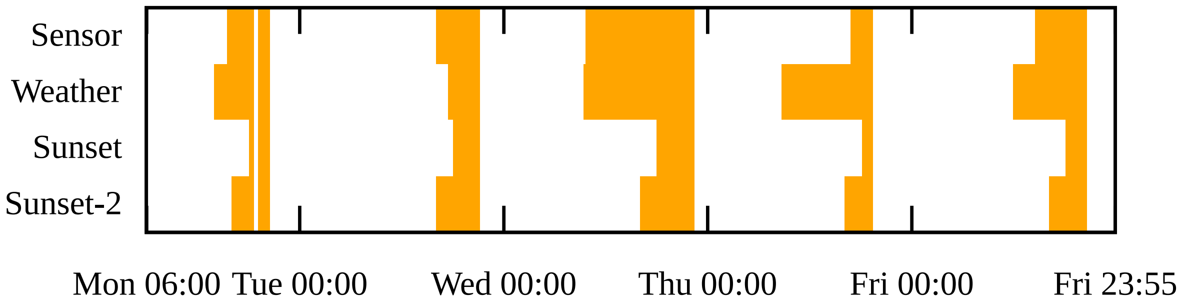

Figure 4.2: Orange shows when lights are on. From top to bottom: ground truth from sensor; when occupied and sun is shining; when occupied between sunset and sunrise; when occupied between 2 hours before sunset and sunrise

RoomMap (Section 2.5.2), it would be reasonable to assume the lights are on whenever the room is occupied.

A weather station on the roof of the building[1] includes a binary sunlight sensor, which indicates whether or not the light level exceeds a predefined threshold. This can be used to guess when the lights are on by assuming that they are whenever the office is occupied and the sun is not shining ('Weather' policy): this was accurate 88% of the time during the week period. Given the limited significance of lighting in the overall energy consumption picture and the fact that this method allows one sensor per room to be replaced by one per building—or even one per city—for a relatively low degradation in accuracy, it seems a reasonable tradeoff. More accuracy could perhaps be obtained by using more detailed weather data and information from OpenRoomMap about the positions of windows in offices.

It is possible to go one step further and attempt to predict lighting state without any sensor data whatsoever, using only calculated sunrise and sunset times. Assuming the lights are on when the office is occupied between dusk and dawn was accurate around 84% of time ('Sunset' policy); refining this estimate to include the two hours before sunset improved its accuracy to around 92% ('Sunset-2' policy, Figure 4.2). Again, knowledge of occupancy is a key piece of data that allows others to be estimated with surprisingly good results and shows that large scale sensor deployments are not always necessary.

4.2 Modelling building energy consumption

Crowd-sourced inventory information can be used even without any additional sensing to produce useful estimates of building energy consumption, disaggregated by function; this can then be used as an input for apportionment. Even where live building-level metering data is available, the ability to estimate a breakdown of *how* this energy is being consumed is valuable for a personal energy meter. This section describes a novel tool for modelling consumption and its evaluation in the William Gates Building. It relies only on data that is easily gathered and sensing which could plausibly be done on a large scale, meaning a personal energy meter could operate even in buildings where no specific energy monitoring hardware installed.

[1] `http://www.cl.cam.ac.uk/research/dtg/weather/`

The model operates by estimating the energy consumption for categories of devices in the building inventory and uses a variety of estimation methods to model different energy use patterns. The summation of energy consumption for each category is then compared with the recorded energy consumption from the building's electricity and gas meters. The category totals can also be used to estimate the breakdown of measured consumption to provide more information to personal energy meter users.

4.2.1 Building modelling tools

There are dozens of existing energy modelling packages [35], though most focus on exploring design options or performing one-off examinations of buildings. Although capable of producing accurate results they require expert users and detailed building survey information. Perhaps the best known example is the DOE-2 software[2] produced by the US Department of Energy. DOE-2 uses hourly weather data to calculate the hour-by-hour performance and response of a building with a known description; heat gains to building spaces are converted to cooling or heating loads on the air using pre-calculated 'weighting factors'. An accessible interface to DOE-2 is provided through the eQuest package.[3] eQuest is designed to make it easier for a single user to capture a building design and parameters in contrast to the new approach described here in which data is reported by a large number of users.

An alternative 'heat balance' method uses a detailed heat model of the thermal transfer processes in the rooms to calculate loads from heat gains; this is generally slower but more accurate. The best known example is the Building Load Analysis and System Thermodynamics (BLAST) system, also supported by the US government. It was developed for predicting energy consumption and systems performance and costs of new or retrofit building designs. A very simplified form of this method is used for the predictions here.

EnergyPlus[4] combines many of the features from these programs. It uses a modular system to permit the construction of detailed building models. Many new building technologies and building and systems simulation models are accessible which represents a significant step forwards in terms of both computational techniques and program structure [34]. TRNSYS[5] is a general simulation package which makes use of modules to model a wide range of systems. The modular and extensible nature of these two systems provides a huge degree of flexibility and both would be candidates for hosting modules implementing the various aspects of the model described here.

4.2.2 Energy estimation methods

A number of energy estimation methods are required to model different device usage characteristics:

[2] `http://www.doe2.com/`
[3] `http://www.doe2.com/equest/`
[4] `http://apps1.eere.energy.gov/buildings/energyplus/`
[5] `http://sel.me.wisc.edu/trnsys/`

4.2.2.1 Simple modulated devices

Constant rate The constant rate method simply assumes a continuous consumption for a device. This is appropriate for always-on devices such as safety lighting, VOIP telephones and printer standby power. The energy consumption of an example of each device is measured using a plug-in power meter and assumed to apply to all devices of the same type.

Timed The building management system in the William Gates Building controls the lighting in public areas (approximately 12 kW) according to a timer. This method applies the measured energy consumption of each device type at a constant rate during the programmed on periods.

Sub-metered It is increasingly common to install sub-metering to monitor the consumption of large energy consumers within a building. In the William Gates Building, for example, sub-meters were used to profile the energy consumption of the machine rooms and associated air-conditioning. These account for an average of 89 kW, which is a significant proportion. Note that sub-metering by building region (e.g. specific corridors) is not directly useful here because many different types of device will be connected to the same circuit. This helps provide a spatial breakdown of energy consumption but not an itemised one. Similarly, the problem of providing an itemised breakdown cannot be fully solved by sub-metering large energy consumers. For example, office lighting for the William Gates Building can consume more than 100 kW (when all switched on). Direct measurement of this would require metering every lighting circuit in the building and then removing the consumption of all the other (timed) lighting from this total.

4.2.2.2 Occupancy-modulated devices

Some devices in the building are switched on and off by occupants and so their power estimate should be modulated by the number of people in the building. This method scales the total power consumption of a set of devices by the proportion of the maximum expected building users currently present, and is applied to devices such as computer monitors and office lighting (see Section 4.1.1 for an evaluation of the accuracy of estimating lighting state based on occupancy).

Here the estimate based on the access logs from the security system described in Section 3.5.3.1 is used. This is a somewhat broad approximation because:

1. the system is based around electronic door locks and so many people can pass through a door for a single access entry; and

2. the system only authenticates ingress events and so only generic unlock events rather than (anonymised) identifiers are recorded when someone exits.

Alternative or complementary approaches might be to determine occupancy based on wireless traffic from smart phones or workstation activity or to use GPS to determine when a user is en-route to or from the building. Better still would be to use a location system as discussed in Chapter 5.

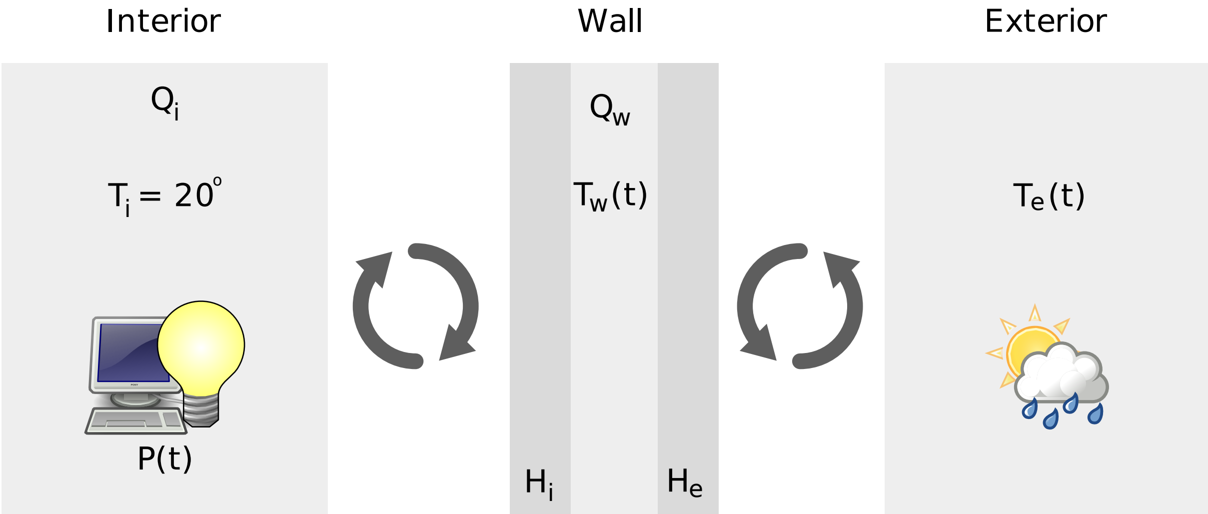

Figure 4.3: Thermal model for the building HVAC [177]

4.2.2.3 Heating, Cooling and Ventilation (HVAC)

The final method provides a simple estimate of the energy consumption of the building's HVAC system. The approach is to estimate the amount of energy required to keep the interior of the building at a desired set-point temperature (T_i) given the heat input to the building from device energy use (including computers) and the heat loss (or gain) due to the outdoor temperature ($T_e(t)$ at time t). The system is modelled as thermal energy movement between three bodies (Figure 4.3):

1. the interior
2. the exterior walls
3. the exterior

The exterior walls act as a buffer between the interior and exterior temperatures. Outdoor temperature is measured by a weather station on the roof of the building.[6] $P(t)$ represents the power lost by devices in the building as heat; each device profile also includes an estimate of its efficiency. The system can therefore be modelled by the following simple differential equations:

$$\frac{dQ_i}{dt} = -H_i(T_i - T_w(t)) + P(t) \tag{4.1}$$

$$\frac{dQ_w}{dt} = C_w \frac{dT_w}{dt} = -H_e(T_w(t) - T_e(t)) + H_i(T_i - T_w(t)) \tag{4.2}$$

The first of these equations gives the HVAC load—the energy required to maintain the internal fixed point temperature T_i (which in this case averages 21 °C). The (numerical) integral of the second equation tracks the temperature of the wall, T_w, over time (t). A negative HVAC load indicates that energy must be put into the building to maintain the temperature (heating demand), and a positive HVAC load indicates that energy must be removed (cooling demand). Thus, raw load can be interpreted in three different ways:

[6] `http://www.cl.cam.ac.uk/research/dtg/weather/`

Wall - solid masonry	2.4 [136]
Wall - modern building standards	0.45-0.6 [136]
Wall - best methods	0.12 [136]
Single glazing	6 [156]
Double glazing	1.4-3.12 [156]
Roof	0.16-0.25 [156]

Figure 4.4: A selection of typical U-values

1. $\left|\frac{dQ_i}{dt}\right|$ is the total power required to maintain the building temperature;
2. $\mathbf{max}(0, \frac{dQ_i}{dt})$ is the total power required to cool the building;
3. $\mathbf{max}(0, -\frac{dQ_i}{dt})$ is the total power required to heat the building.

H_i and H_e correspond to the 'leakiness' of the wall towards the interior and exterior of the building respectively. These are derived by multiplying the surface area (m^2) by the thermal transmittance or U-value ($W/m^2/K$). The surface area of the building was estimated manually (40,000 m^2) but such information could also be obtained from the floor area in OpenRoomMap and an estimate of ceiling height and roof-pitch. U-values are normally quoted for a single surface and a typical value suggested by MacKay for best building methods of 0.15 was adopted [136]—the building in question won an architectural award for its heating and cooling efficiency.[7] The model uses separate U-values for the inner shell (U_i) of the wall and the outer shell (U_e) and so it is further assumed that the outer shell has 2.5 times the thermal resistance of the inner shell. Given that U-values combine in the same manner as resistors in parallel:

$$\frac{1}{U} = \frac{1}{U_i} + \frac{1}{U_e} \tag{4.3}$$

and substituting $U_e = 2.5U_i$ gives:

$$U_i = 3.5U = 0.53 \tag{4.4}$$

$$U_e = \frac{3.5U}{2.5} = 0.21 \tag{4.5}$$

Although this model is very simple, it does produce acceptable results (see Section 4.2.3) and so serves the purpose of demonstrating that little input data is genuinely necessary; of course, it could be replaced with more sophisticated physics in subsequent versions.

4.2.3 Results

Figure 4.5 shows the model output and recorded electricity consumption in kW for November 2009 to August 2010. The dark 'metered consumption' line is the half-hourly measurements of electricity consumption as recorded by the electricity company. The categories in the breakdown are as follows:

[7] `http://www.cabe.org.uk/case-studies/william-gates-building/design`

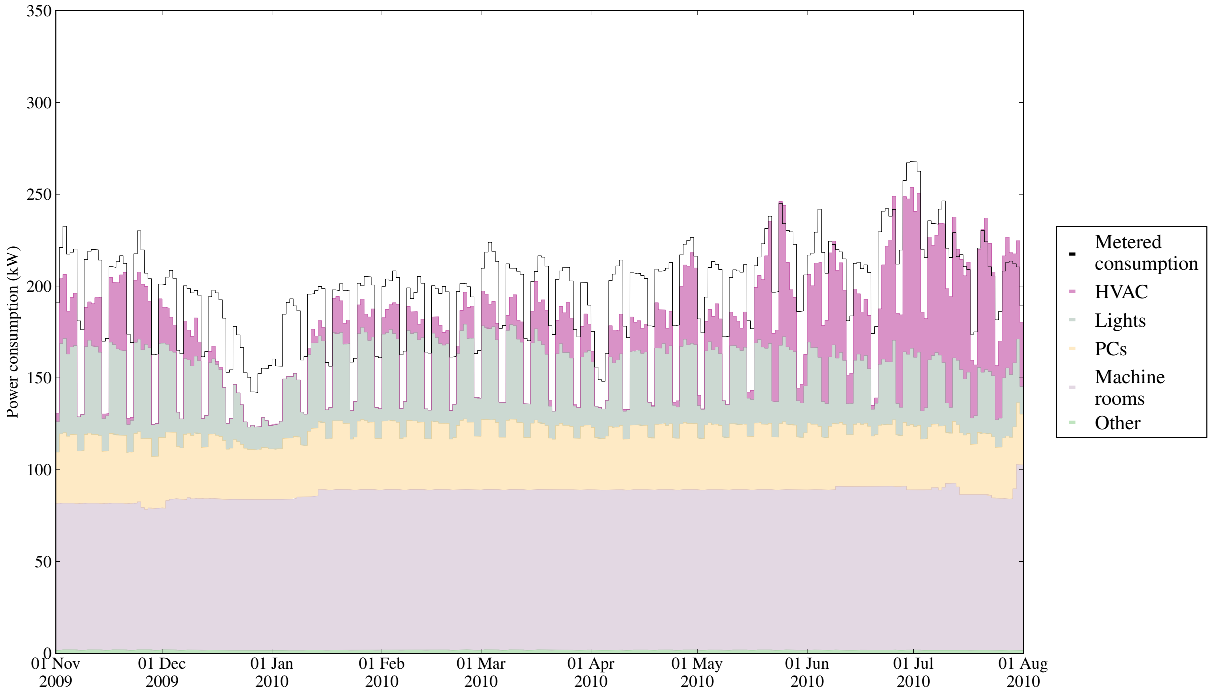

Figure 4.5: Daily breakdown (Nov 09 to Aug 10) shows trends in electricity consumption are correctly estimated [177]

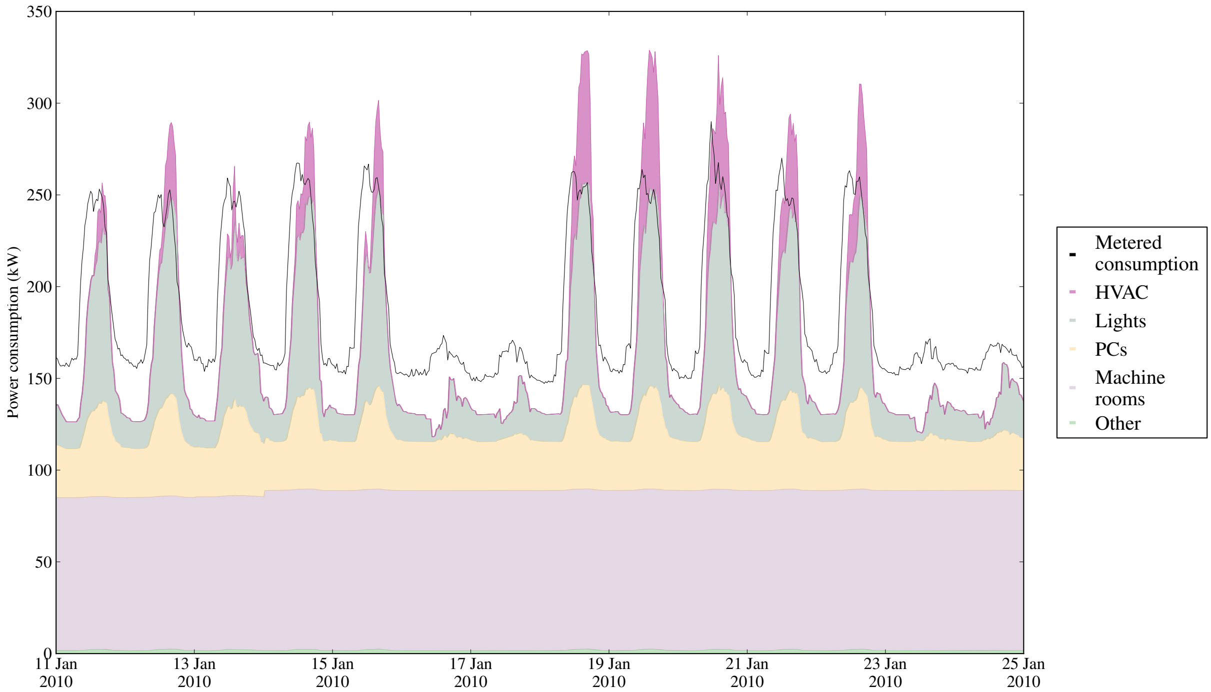

Figure 4.6: Half-hourly breakdown (Jan 2010): electricity requirements during winter vary mostly due to lighting needs [177]

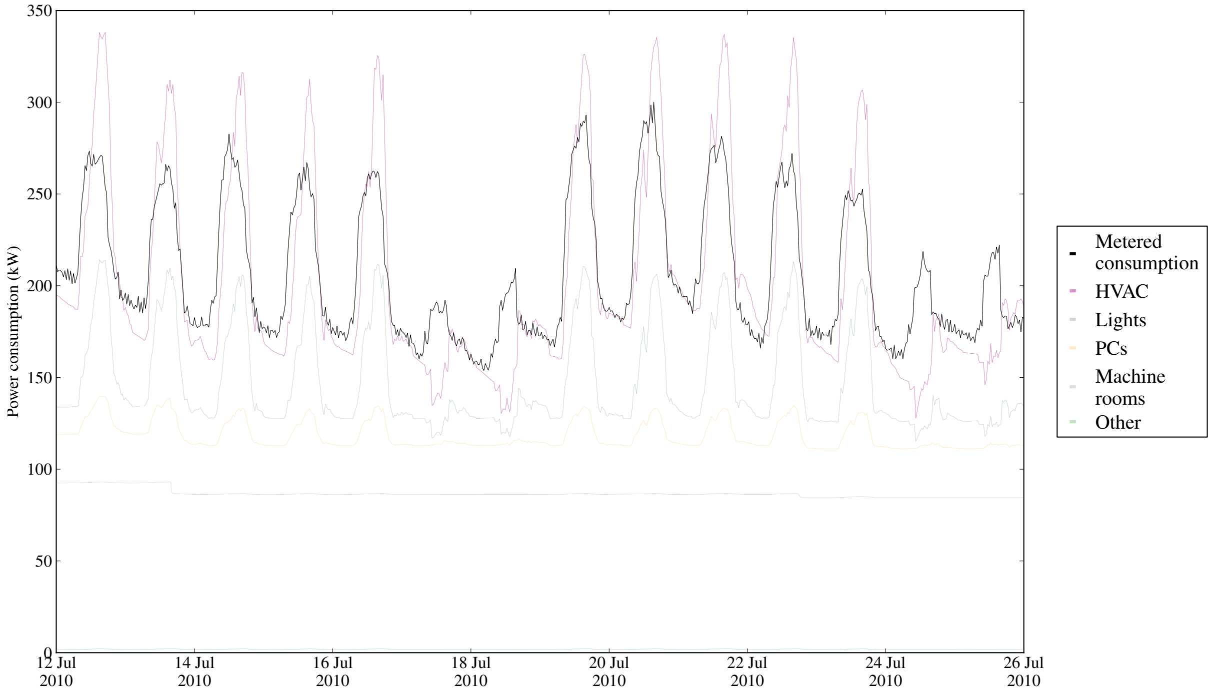

Figure 4.7: Half-hourly breakdown (Jul 2010): cooling dominates the electricity requirements during summer [177]

HVAC is the output of the heat model for the building. Initially only cooling is considered to account for electricity usage.

Lights includes lighting within offices (modulated according to the occupancy of the building) and in public areas (modulated according to a timer function).

PCs covers the energy use of personal computers and monitors in offices, assuming that the PC itself is left on continuously whereas the monitors are switched on or off according to the occupancy of the building. Both are assumed to consume 70 W.

Machine rooms considers servers, uninterruptible power supplies and air conditioning units in the machine rooms. This is a mixture of sub-metered readings and manual estimates.

Other contains minor items from the OpenRoomMap inventory such as printer idle power, telephones and a small number of electric heaters.

Notable from the graph is that the predicted consumption displays similar trends to the true measured value. Over the annual period the load on the HVAC system increases during the summer months and falls to nothing over the Christmas period when the building is quiet and the exterior temperature is low. Unfortunately, no ground truth data disaggregated by function could be obtained as the equipment in each category is distributed throughout the building so a very large number of sensors would be required. Nevertheless, the HVAC estimate fits the trends in the metered electricity consumption particularly well in the summer months when the heating load is highest, suggesting that it is indeed responsible for many of these variations.

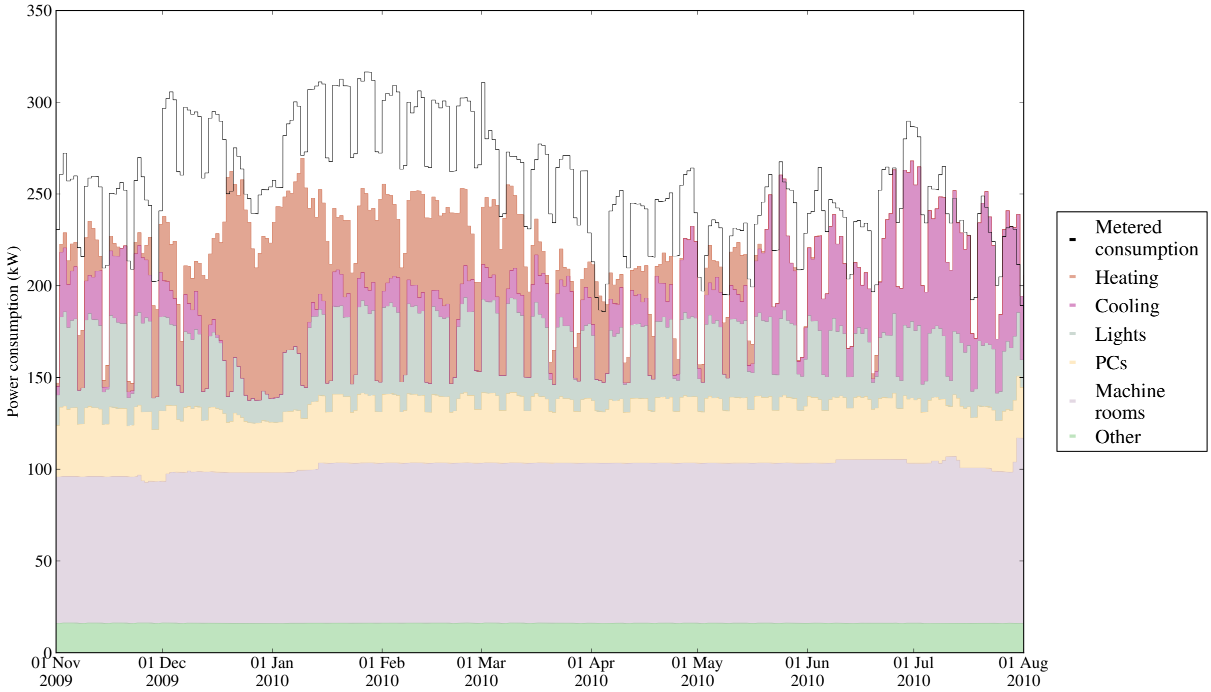

Figure 4.8: The model underestimates combined heating and cooling energy consumption during winter. Note that metered consumption here includes both electricity (recorded half-hourly) and gas (recorded only monthly and interpolated) [177]

Figure 4.6 shows a two week period in January. The peaks in consumption during working weekdays are clear in the model and the breakdown shows that this is mostly due to lights being switched on (in offices). Figure 4.7 shows a two week period in July. In this case the HVAC energy usage is significantly higher due to higher outdoor temperatures.

The effect of including heating in the model is now considered. The heating system is assumed to be 70% efficient and from the gas consumption over the summer when no space-heating is needed an additional cost of 1.4 kW for water heating is derived which is included in the 'Other' category. There is what seems to be a more significant deviation from the measured trace (Figure 4.8, showing both electricity and gas[8]). However, this is due in part to the fact that the gas consumption data for the building is measured monthly and must therefore be interpolated linearly so day-scale changes in consumption as predicted by the HVAC model are not reflected in the measured consumption trace (Figure 4.9). There are many factors which could be altered to obtain a better fit, such as changing the U-value of the building, the efficiency of the heating system or the fixed point temperature but this is not done for fear of over-fitting what is a very simple model. The results are sufficient to show how a personal energy meter could provide individuals with useful insights into the breakdown of their consumption without the need for extensive device-level metering.

[8]Converted to kWh in accordance with *Conversion factors - Energy and carbon conversions - 2010 update (CTL113)* `http://www.carbontrust.co.uk/publications/pages/publicationdetail.aspx?id=CTL113`

4.2.4 Energy saving scenarios

One aim of the personal energy meter is to suggest savings which could be made, and an advantage of modelling, even if live metering is also deployed, is that it can also be used to consider the effects of hypothetical energy-saving scenarios. The model for the William Gates Building suggests three big potential areas: machine rooms, PCs and lighting. Figures 4.10 and 4.11 show the results of these scenarios. Of course, reducing computing load also reduces the heat output from these computers and so decreases the cooling required in summer (leading to further savings) but increases the heating required in winter; the integrated model is therefore more valuable than simple calculations as it takes these effects into account.

Normal computing An estimate of the energy consumption of the building if (like many other buildings on the estate) it contained no significant server infrastructure and a single workstation per occupant.

PCs off An estimate of the impact of building occupants switching off all workstations when not in the building.

LED lighting An estimate of the impact of switching to LED lighting, replacing the current 50 lm/W lighting with LED equivalents achieving 160 lm/W.[9]

Finally, since the goal is to produce a modelling tool which can be automatically applied across many buildings the sensitivity of the model to the building U-value should be considered. Figure 4.12 shows the result of running the simulation with 4 different U-values. These results show that a good choice of value is probably around 0.15. Its clear that the wrong choice of U-value can have a significant impact on the quality of fit. However, it is easy to notice that the fit is incorrect. If building-level metering is available, one technique might be to collect data as to the point in the year when the building's heating system is first switched on for a significant period of time and to adjust the building U-value to produce a similar effect.

The model can be used by a personal energy meter to estimate what proportion of the consumption belongs in each category of device, making the feedback presented to end users more informative. This strategy is adopted in the model described in Section 6.4. If the inventory is believed to be representative then consumption left unaccounted for can be attributed to errors in the model, and it is reasonable to scale the estimates for each category up so their total matches the true consumption; alternatively, the remainder can simply be apportioned in the same way that the entire building's total would have to be if no model existed.

4.3 The need for fine-grained measurements

The estimates made by the model described in the previous section can be incrementally improved by adding sensing; more devices can be metered directly and their actual consumption fed back into the model to help calculate the heating or cooling load as well as

[9]The US Department of Energy estimates that 160 lm/W LED lighting will be market-ready by 2025: `http://www1.eere.energy.gov/buildings/ssl/efficacy.html`

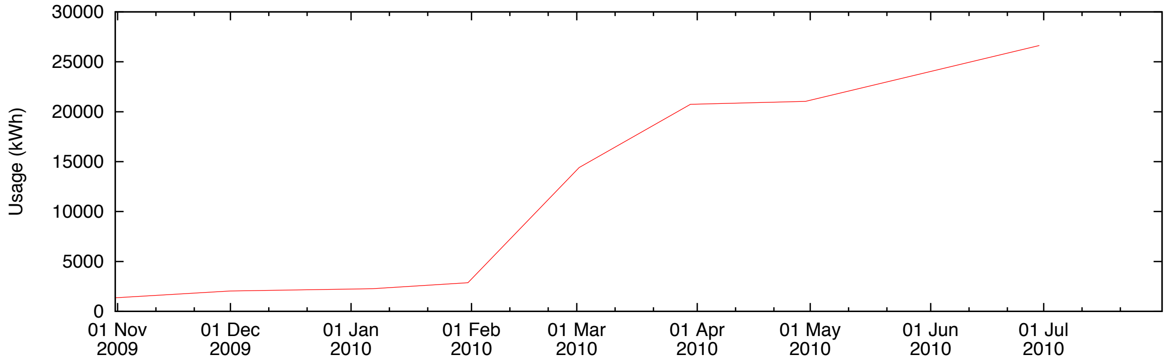

Figure 4.9: Gas consumption is recorded only monthly

Scenario	Av. Power	Change	Saving
Metered	275 kW		
Current conditions	213 kW		
Normal comp.	118 kW	95 kW	£83,000
PCs off	206 kW	7 kW	£6,100
LED lighting	192 kW	21 kW	£18,000

Figure 4.10: Predicted reductions in average power consumption over a year

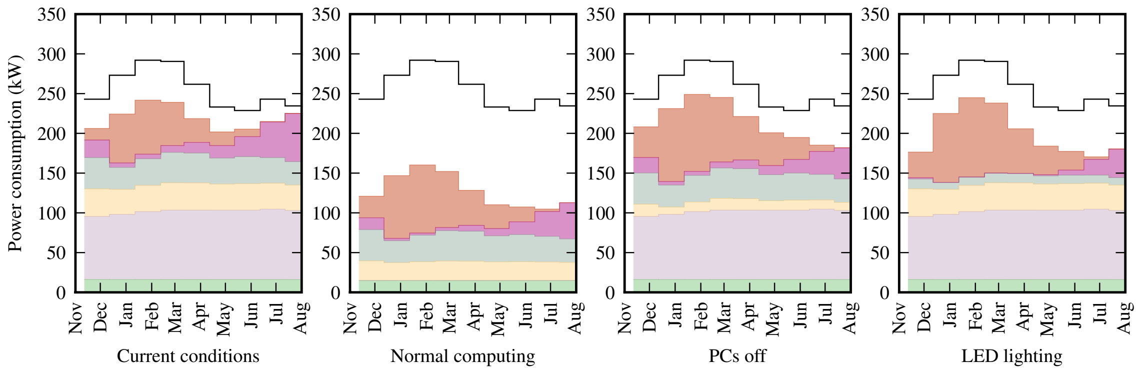

Figure 4.11: Increases in heating load and decreases in cooling load follow from energy savings [177]

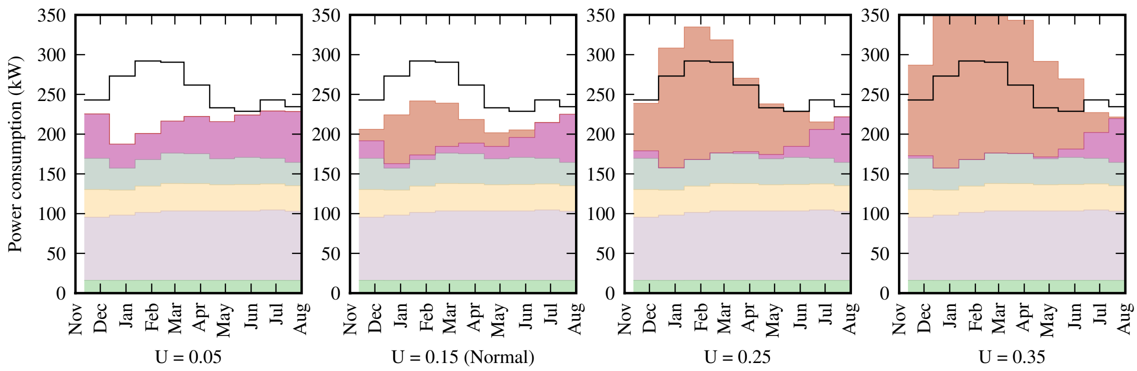

Figure 4.12: Varying the choice of U-value has a significant impact on the model prediction [177]

being reported to a personal energy metering system. This works well for owned resources as described in Section 3.6.1, whose entire consumption can be allocated to their owner, but as soon as a device is shared, knowing its total cumulative energy consumption is not enough: it is necessary to understand *how* and *why* that energy was used in order to determine who was responsible. In the simplest cases of devices which consume energy only when in active use, knowing the instantaneous consumption may be sufficient since the entire consumption may be allocated to the person using it at the time. However, many devices can serve more than one user at a time: even simple things like printers and coffee machines use power to heat fuser elements or water that benefit all subsequent users. It is therefore necessary to understand the energy cost of printing a page, or making a cappuccino, and separate it from the base costs.

Computers are also an excellent, and significant, example of this class of device. By 2007 Gartner was estimating that datacentres accounted for almost a quarter of global CO_2 emissions attributable to IT and by themselves as much CO_2 as the aviation industry.[10]

The world has moved away from the mainframe model of hundreds of users sharing a single machine through to the era of the personal computer to the dawning age of Weiser's ubicomp vision [202]. The previous sections have discussed the problem of the proliferation of electronic devices whose consumption must be metered; however, with the growth in popularity of cloud computing the model is returning full circle, with computing power centralised and shared amongst, in some cases, millions of users. Every user may have dozens of devices but expects to access the same data and applications on each; the natural tendency is therefore towards applications as services, potentially hosted by third parties and accessed over the Internet. This idea has been discussed for a number of years, originally in the context of thin clients and remote desktop systems such as VNC [180]; it is the explosion of web applications that now makes it seem an attractive proposition. Clearly, some users will be much more active than others, with some people merely registering once to explore a site and never visiting again while others rely on it for their daily business. In order to allocate a fair share of the energy costs of shared devices to each individual based on his actual consumption, it is necessary to meter power consumption at a far finer level than discussed so far, investigating the energy implications of each action a user commands.

4.4 Manual device profiling

As discussed in Section 4.1, understanding the power states of a device and the energy costs of actions it performs allows the total cost of an individual's usage pattern to be calculated. The simplest way to build a device profile is to study the power trace of a sequence of known actions. This section demonstrates this technique for two different devices.

As a middle ground short of continuous online measurement, custom hardware was built based initially on the design used by Hylick et al. to analyse hard drive energy consumption [101]. An ATMEL microcontroller was used to integrate the readings of an off-the-shelf clamp meter. Figure B.2 on page 181 shows the schematic circuit diagram; a current probe designed for microscopes[11] was attached to and powered by the board

[10] `http://www.gartner.com/it/page.jsp?id=530912`
[11] A LEM PR 20

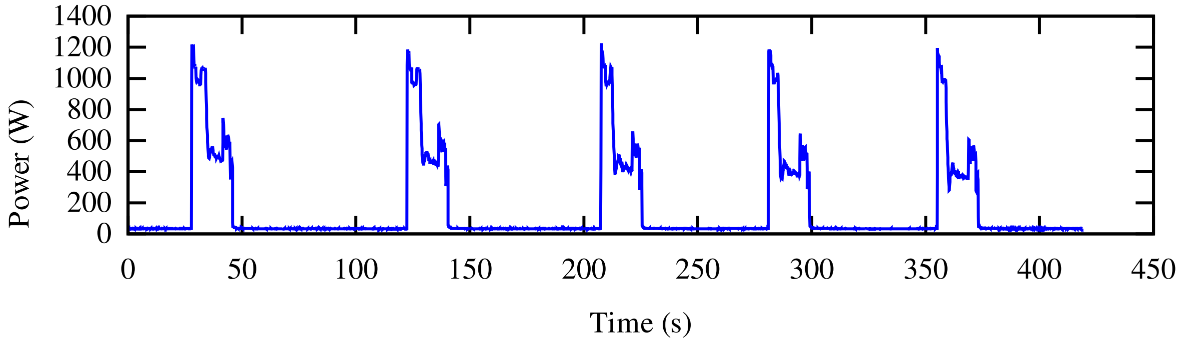

Figure 4.13: Power drawn printing five single pages

and the timestamped results were logged to a memory card at 20 Hz. This frequency allows the resulting trace to be aligned by hand with an activity log so the energy costs of specific events, such as printing a page or making a cup of coffee, can be identified.

4.4.1 Coffee machine

As one example, the consumption of a shared coffee machine[12] was measured using the apparatus described above, which revealed that the additional electrical energy required to make a single cup of coffee is approximately 62 KJ, or 0.02 kWh. The standby consumption is 4 W.

4.4.2 Printer

The same technique was used to measure the energy consumption of one printer.[13] An example trace is shown in Figure 4.13. The printer draws 32 W when idle (in power save mode), and consumes, on average, an additional 11,200 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; the average energy cost per page for the whole workload over several days was 8,720 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.

The energy measurement can be combined with the logs from the print server to improve on the 'personal load' policy described in the apportionment case study in Section 3.5 (Figure 3.7). The energy consumed by a particular print job was originally spread over all occupants and so for each job a share of the energy consumed is removed from each person's allocation before the total is reassigned to the individual who printed the material. The results of this policy for one staff member, who printed a large set of lecture course material in the week in question, are shown in Figure 4.14 and represent an increase of 8 kWh for the week. Neither of the other example users referred to in Section 3.5 printed anything; their results show a reduction in allocated energy of around 0.5 kWh for the week.

[12] A Jura IMPRESSA X7
[13] A HP LaserJet 4200dtn

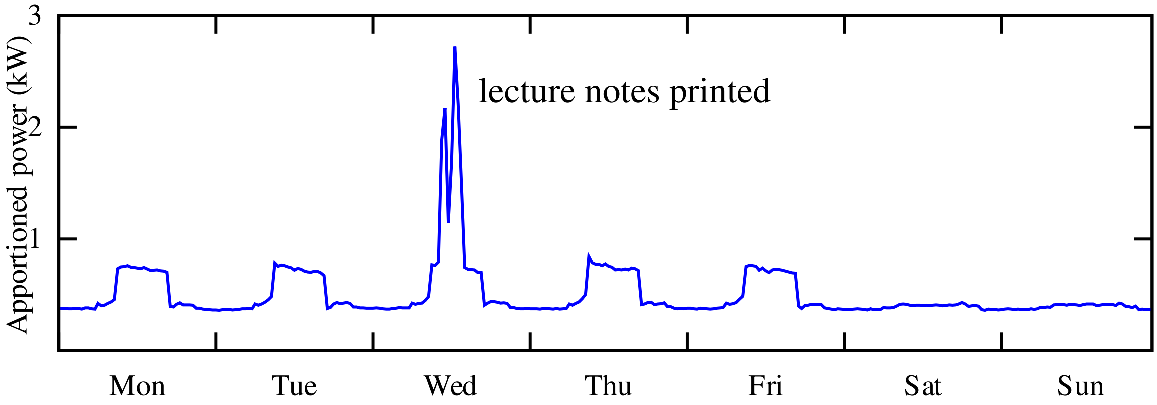

Figure 4.14: Apportionment with printing costs

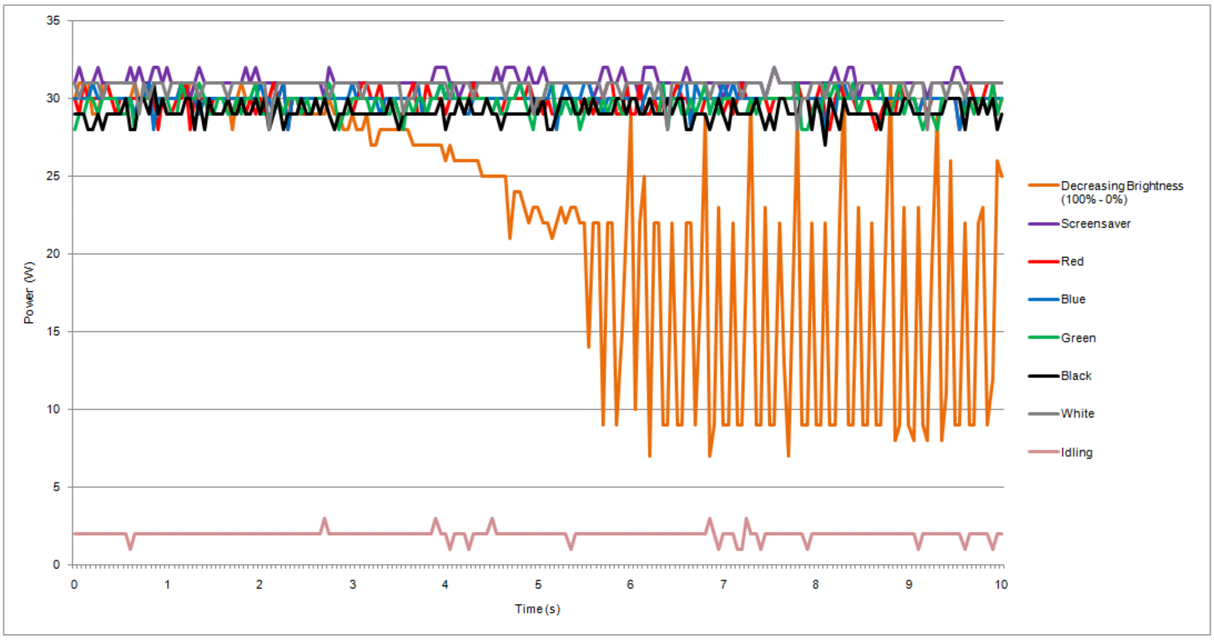

Figure 4.15: Energy consumption of an LCD monitor with different settings and displays

4.4.3 LCD monitor

Profiling a LCD monitor[14] showed it has an average idle consumption of 2.0 W and average active consumption of 30.0 W, but the high frequency of measurements revealed that the power draw is in fact dependent on the brightness setting and the colours and movement displayed. At full brightness, a black screen used 29.2 W on average while a white one used 30.8 W; displaying a moving screensaver used 31.1 W, probably because of the additional switching required. Figure 4.15 shows the full results, and also shows that reducing the brightness also reduces the energy consumption approximately linearly to an average of 16.4 W. Despite these minor variations, the data supports the simple profile used in the building model.

[14] A Samsung 170T

4.5 Automated profiling of programmable devices

Determining the energy costs of tasks on devices as described above is a manual and lengthy process; it becomes vastly more complicated on programmable devices because of their much greater range of possible activities. Furthermore, these devices tend to perform operations much more quickly—millions of times per second, in the case of microprocessors—making manual alignment of power traces with activity logs very difficult. This section therefore outlines a mechanism to achieve task-level metering for a programmable device, allowing a profile to be built up of the energy costs at a function call level based on taking a large number of measurements of particular operations in an automated fashion. Automating the process means each measurement can be repeated to reduce error.

4.5.1 Requirements

There are a number of desirable properties that an automated framework for decomposing the power measurements of devices should exhibit to ease the task of building profiles:

Automated test execution The tests should proceed (as much as permitted by the device) without requiring user interaction. This removes a major source of variability and allows the tester to increase confidence in a result by repeatedly executing the same sequence of actions to eliminate random noise (but not systematic error).

Batch operation It should be possible to run a whole sequence of tests without intervention.

Untethered operation No physical connections (except those for the metering itself) should be required required to any of the interfaces on the test device.

No hardware modification It should not be not necessary to modify the test device at all other than to fit a meter and install a standard application.

4.5.2 Implementation

To achieve these goals, a novel power measurement system was devised. The system is centrally orchestrated by the Power Server, which is responsible for sending test scripts to the device to be profiled and collecting and aligning the various traces and log files. A single client program is run on the test device itself which is responsible for acquiring a test script and executing the required actions. Test scripts written in a purpose-designed simple language are interpreted dynamically without requiring any changes to the software running on the test device. The client program collects a timing log of these events which is uploaded to the Power Server at the end of the test. There are no networking requirements beyond the initial download of the test script and final upload of the results. This means, for example, that tests can easily be run examining the costs of changing network or using an external device.

The client running on the test device process proceeds as in Figure 4.16. It first connects to the Power Server and downloads a test script. It then enters the preparation phase

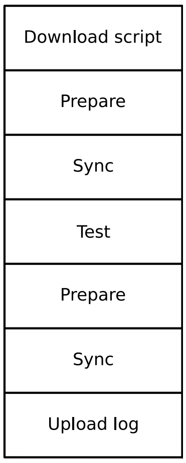

Figure 4.16: Execution stages

and stabilises its power consumption. A predetermined sequences of actions is performed to create a synchronisation pulse. This is used in the data analysis phase to correlate the timing log from the device with the recorded data. The client next executes the test script recording the time at which each action is performed. Once the script is complete the power consumption is stabilised once more and a final synchronisation pulse is emitted before uploading the timing log back to the Power Server. These phases are described in more detail in the remainder of this section.

4.5.2.1 Measurement hardware

The system for decomposing measurements is independent of the mechanism for metering consumption; different techniques will be necessary for different types of device. The examples described in this dissertation were obtained using either versions of the custom measurement hardware described in Section 4.4 or a commercial sampling board, but many of the research or off-the-shelf metering devices described in Section 2.3.1.3 such as the ACme [104] would also be appropriate. The main constraint is the sampling rate; the more frequent the voltage and current measurements, the more detailed the resulting profile will be.

Although modern motherboard designs often contain sophisticated mechanisms for reporting power consumption to the operating system these tend not to be suitable for profiling purposes since sampling at the necessary rate itself entails a significant additional power draw. If the metering is performed on a separate device this source of potential inaccuracy is removed.

4.5.2.2 Stabilising the power trace

The primary goal is to break down the overall power consumption of the test device into its constituent parts. Ideally one would identify components which are causing variation in the trace, characterise their consumption and then switch them off. For components such as the CPU this is not an option because the operating system is preemptively multitasking and so other processes are intermittently waking up and consuming resources:

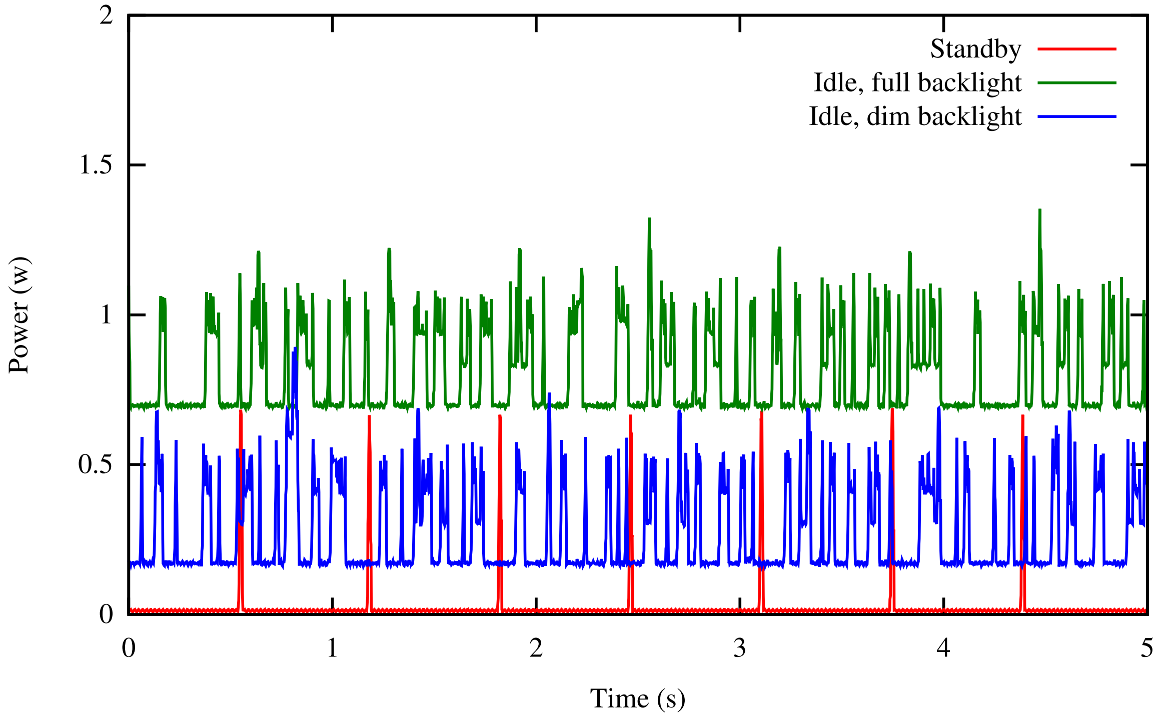

Figure 4.17: Energy consumption of a G1 mobile phone when idle [176]

Figure 4.17 shows the variation in power when even an example mobile phone[15] is ostensibly idle. Instead, a low-priority background process runs in a busy-loop. This consumes all spare CPU cycles and contributes greatly to stabilising the power trace. A small uncertainty is introduced by this technique because it is not possible to distinguish between CPU load created by the test and the background load. However, for the purposes of understanding the peripheral hardware in the device (such as the networking hardware) this should have little effect.

4.5.2.3 Trace synchronisation

Many of the features in the energy traces of electronic devices last only a fraction of a second. For example, a scan for available wireless networks lasts around 500 ms, while the transmission of a single packet takes only a few milliseconds. For this reason it is important to align precisely the times recorded for different events on the test device with the samples recorded on the measurement PC. This alignment allows annotation of the power trace at each instant with the action taking place on the device.

This is achieved by embedding a synchronisation pulse *inside the energy trace* by switching on and off a component with significant power draw in a predetermined pattern. On a mobile phone, for instance, switching the backlight from off to full brightness increases the power consumption of the device by more than half a watt over a period of a few milliseconds (Figure 4.17); on a desktop PC, varying the CPU load can make a difference of tens of watts (Figure 4.18). This effect can be exploited to embed two easily-recognisable 32-bit Gold codes [69] at either end of the trace, with on representing a 1 and off a 0. The pulse sequence is shown in Figure 4.19.

[15]`http://www.htc.com/www/product/g1/overview.html`

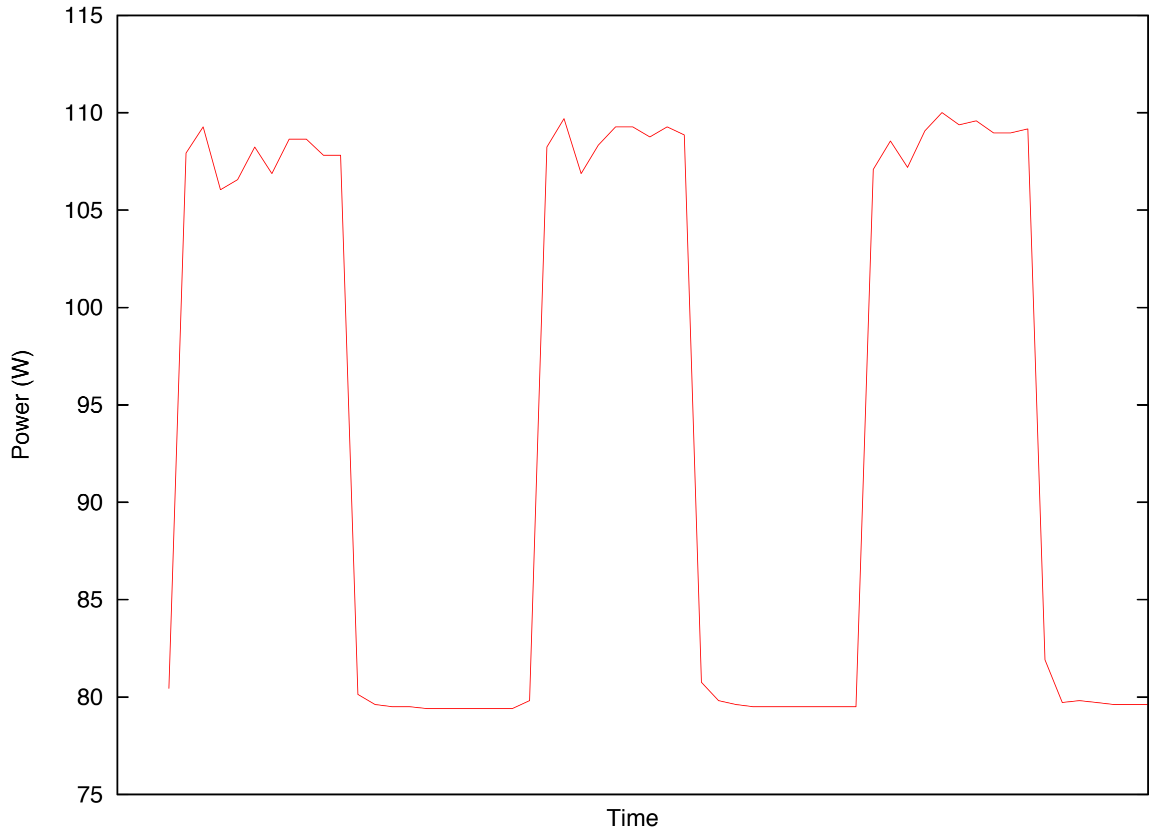

Figure 4.18: Part of a synchronisation pulse produced by varying the CPU load on a desktop PC.

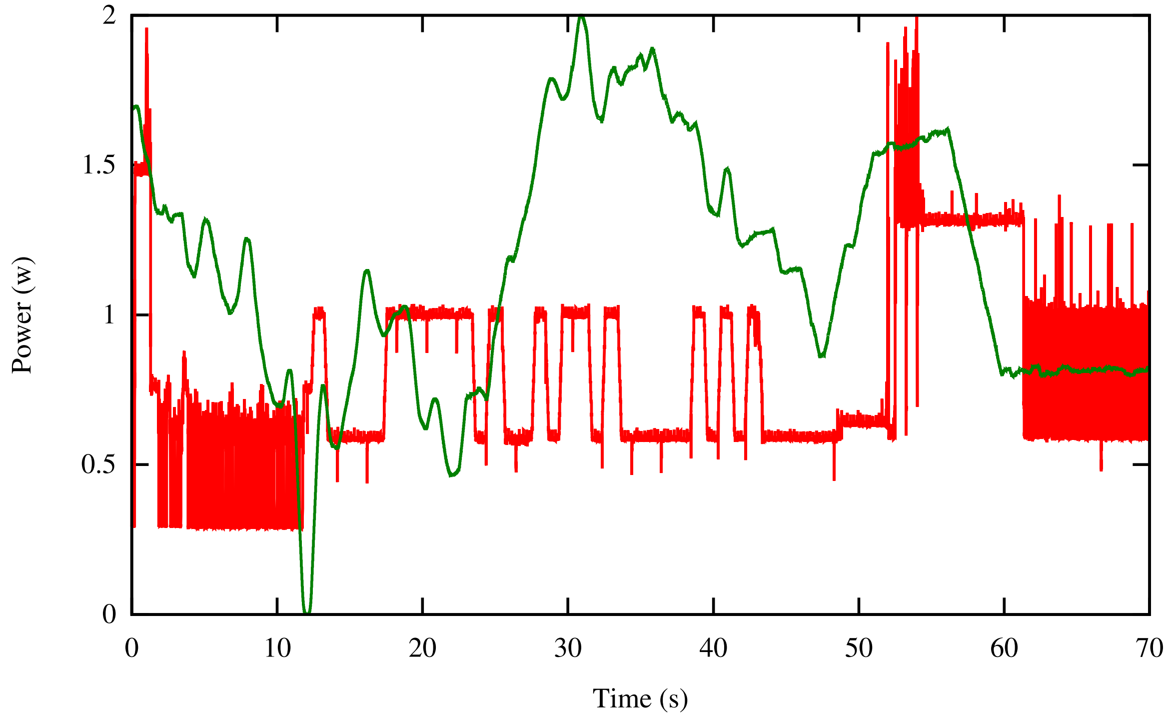

Figure 4.19: A synchronisation pulse (in red, approximately between seconds 12 and 42) and the SSD function with the hypothesised signal (in green) [176]

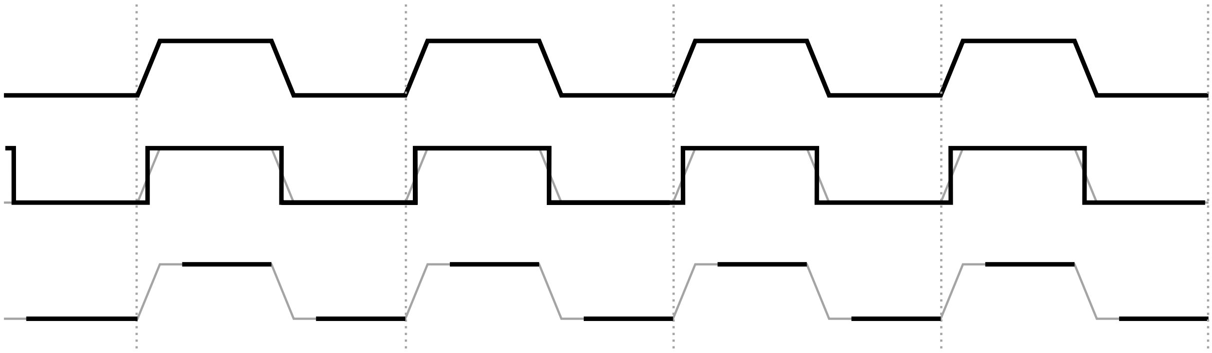

Figure 4.20: A partial hypothesis trace (bottom) is necessary because a square hypothesis (middle) will misalign on the true trace (top) [176]

At the end of the test the timing log recorded by the device contains the switching times for each edge in the synchronisation pulses. These times are combined with estimated values for the power consumption in the two backlight states to generate a hypothesised sequence representing the expected power trace for the pulse. It is valid to assume in this case that the relative drift of the two clocks over the synchronisation time period will be negligible.

To find the sample which corresponds to the start of the synchronisation pulse, a function of the power trace and the hypothesised signal is computed by incrementally increasing the offset of the latter and computing the sum of the squares of the difference (SSD) between the signals (also shown in Figure 4.19). This relatively simple function works well as a measure of similarity and drops to zero at the point where the two signals line up; the fact that the Gold codes have small cross-correlations within a set ensures it is very unlikely that they will match in the wrong position. In all tests a visual inspection has shown this to be a robust way to determine the sample number of the start of both pulses. The cross-correlation function could equally have been used instead of the sum of the squares of the difference, and any pseudorandom bit sequence could have been used in place of Gold codes: the Gold codes simply minimise the chance of the pattern repeating and so matching well against an offset version of itself.

Components take time to switch on and off, and this problem is sometimes exacerbated by software. For example, some recent versions of operating systems fade the display backlight on and off as a visual effect, resulting in diagonal lines in the synchronisation pulses on the power trace rather than sharp edges (Figure 4.20(top)). Attempting to match the square hypothesised signal against this results in them lining up incorrectly, with the vertical edges half way along the sloping lines where we should have the rising edge at the base of the upward sloping line and the falling edge at the top of the downward one (Figure 4.20(middle)). To counteract this, the edges must be removed from the hypothesised signal to leave gaps where the slopes will be.

Changing hardware power state generally requires an inter-process or kernel call so there is sometimes a slight delay between making the API call and the change taking effect. The start of each hypothesised pulse section is therefore also left blank (Figure 4.20(bottom)). These blank 'don't care' section mean that there are multiple points along the trace that are good matches, so the latest maximum is chosen.

This calibration procedure gives four values: s_1 and s_2 corresponding to the sample number for the start of the first and second pulse, and t_1 and t_2 which correspond to the time that the first and second pulses began. From this the sample rate between the two pulses, r can be calculated:

$$r = \frac{s_2 - s_1}{t_2 - t_1} \tag{4.6}$$
$$o = t_1 - s_1/r \tag{4.7}$$

where o is the time offset between the start of the power log and the start of the timing log. The sample s which corresponds to some time t is therefore calculated as

$$s = r(t - o) \tag{4.8}$$

Once the alignment has been calculated in this way, new estimates are formed for the power consumption with the component on and off and the calibration process is re-run with a new hypothesised signal based on these estimates to ensure as accurate a result as possible regardless of the physical device.

4.5.2.4 Baseline calibration

Many of the tests involve switching on some component of the test device, waiting for it to initialise and then examining the additional energy costs of using the device. This process is supported by embedding baseline power calibration in the test scripts. The test writer first annotates the script to indicate that a particular set of steps should be used as calibration information. When processing the log files the average power consumption of these steps is computed and subtracted from subsequent steps. As an example, this allows the energy cost of transmitting data over a network interface to be dissociated from the ongoing power requirement to keep the interface active.

The example script in Figure 4.21 measures the power consumption for switching on the Wireless LAN, holding it on for baseline calibration and sending 25 KB of data using a 1448 byte buffer. Each line of the script corresponds to a step in the test and contains a number of fields separated by the ':' character. The first field indicates the type of action to be taken by the measurement framework. The most relevant of these are the `BASELINE` action which informs the measurement system to calculate a baseline power consumption using the average power consumption for the duration of the step. The `MEASURE` action causes the power consumption (minus the current baseline power measurement) of the step to be recorded. The average power consumption (watts) and the total energy consumed (joules) are calculated. The number in the second field indicates the number of units of activity present in the step which is used to produce a normalised, unit-cost, energy consumption. In the example there are 25 units in the `SendTCP` step sending 25 KB of data. This causes the system to automatically calculate the cost of sending 1 KB of data and add it to the logs. The `MEASURE_CONT` action indicates to the system that it should treat this step as a continuation of the previous one and make a measurement over the entire duration of both.

```
NONE:1:ToggleWakeLock:true                  # Force the device to remain awake
NONE:1:ToggleTelephony:false                # Disable the cellular radio
NONE:1:ToggleWifi:false                     # Disable the WiFi radio
NONE:1:WaitTelephonyDisconnect
NONE:1:WaitWifiDisconnect
NONE:1:ToggleCPU:true                       # Start a background busy thread
PRESYNC:1:SetBacklight:1                    # First synchronisation pulse
NONE:1:DoSleep:800
PRESYNC:1:SetBacklight:0.1
...                                         # Remainder of sync pulse omitted
BASELINE:1:DoSleep:5000                     # Calibrate baseline power
MEASURE:1:ToggleWifi:true                   # Enable the wireless network
MEASURE_CONT:1:WaitWifiConnect              # Wait for a connection
NONE:1:DoSleep:5000                         # Wait for 5 seconds
BASELINE:1:DoSleep:10000                    # Calibrate WiFi idle power
MEASURE:1:OpenSocket:192.168.0.210:8060:1   # Open a TCP connection
NONE:1:DoSleep:5000                         # Wait for 5 seconds
MEASURE:25:SendTCP:25:1448:false            # Send 25 KB TCP using 1448 byte buffer
MEASURE_CONT:1:DoSleep:3000                 # Wait for 3 seconds
NONE:1:CloseSocket                          # Close the TCP connection
...                                         # Start of sync pulse omitted
POSTSYNC:1:SetBacklight:0.1                 # Trailing synchronisation pulse
NONE:1:DoSleep:800
NONE:1:ToggleCPU:false                      # Release CPU
```

Figure 4.21: Parts of an example test script

4.5.3 Results

The technique can be applied to anything from a printer to a rack of servers, provided it is possible to write an interpreter for the scripting language; a case study using Android-based mobile phones as a testbed is described in Appendix A and the power measurements described in Section 5.7 demonstrate its utility and flexibility.

One measure of the accuracy of the synchronisation is to look at the variance in r across runs in the testbed deployment described in Appendix A. In a perfect world this would have a value of 4,000 and the experimental results were never out by more than 3 ns.

4.6 Summary

Understanding how energy is being used is an important first step to improving efficiency. It is impractical to sense directly the consumption of each energy consumer within a building; as an alternative, it can be estimated based on inventories, device profiles and second-order indicators of use such as occupancy.

This chapter has described a modelling technique which could be practically applied across many buildings. The system described can operate with a minimal amount of live sensing information but could still extend to accommodate more sources as they are installed. Minimising the effort involved in initial data collection is important to this goal and so OpenRoomMap is used to crowd-source this information from building users.

To create device profiles one must break down the consumption of a device into the energy costs of its various actions. This chapter has shown how this can be achieved manually using custom hardware; it has also described a measurement framework which can be used to create a fine-grained understanding of the energy consumption of a programmable device. The synchronisation information required can be embedded in the measurement trace itself, making the entire process automated and repeatable.

Chapter 5

Pragmatic location tracking

Contents

Overview

This chapter demonstrates that location can be used to determine the usage patterns of shared resources, and to apportion energy costs to individuals. It discusses the problems with today's location systems and identifies the key characteristics required of a system for use in a personal energy meter: *low cost*, *low infrastructure* and *easily deployable*. The most successful approaches to date have repurposed existing infrastructure and user devices to provide tracking. Bluetooth has been a particularly popular tracking medium due to its ubiquitous implementation in modern devices, its low power design and its low cost componentry. This chapter enumerates the different techniques on which Bluetooth tracking can be based, some of which are novel, evaluates their properties experimentally and theoretically and shows how to use them to construct a large scale tracking system suitable for apportioning energy use.

Some of the contributions presented in this chapter have also appeared in a separate publication [87]. Figure 5.3 is reproduced courtesy of Robert Harle.

5.1 Context awareness for personal energy metering

A research group in the William Gates Building shares a coffee machine. In this case, as for most shared resources, there is no second order information available on its usage. All that are left are several possible options that require varying investments of time and infrastructure. Devices like this could be counted as part of the base load and their usage ignored entirely; users could be asked to keep track of their usage manually; some form of identity prompt could be deployed, such as the PINs often required on photocopiers for accounting purposes, or a separate sensor system could be installed. The most appropriate will depend on the significance of the resource in question.

Measurement apparatus described in Section 4.4 showed that two cups of coffee per day accounts for about 3% of an individual's 150 W personal load (see Section 4.4.1 for further details). This is insignificant relative to the power draw of the whole building and it may not even be worth the energy cost of a sensor system to apportion its use. Even ignoring it altogether would probably be justifiable—but lessons learnt from the coffee machine can be applied to all sorts of other equipment, so it is a valuable case study.[1]

To evaluate the feasibility of the manual method, during the course of one week members of the research group were asked to make a mark against their name on a tally sheet every time they had a coffee. An 'Anonymous' row was also included 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. An equal apportionment policy would therefore divide the total energy cost (20,646 kJ) amongst everyone and allocate 390 kJ each for this week. If instead the energy costs of those cups of coffee that were logged were allocated to the appropriate individuals and then divided the costs of the remainder equally amongst everyone, the mean difference from the equal policy is 63%. Finally, assuming that the logging method was entirely accurate and captured all cups of coffee made and therefore all energy costs are allocated to the individuals responsible, the mean difference from the equal policy is 164%.

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. Any attempt to apportion the use of these resources as part of a 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 an 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.

Contextual information [43, 46] will therefore provide crucial indications for apportioning the use and energy costs of resources. This context can come from many different sources;

[1] The University of Cambridge Computer Laboratory has a long and proud tradition of augmenting coffee machines (`http://www.cl.cam.ac.uk/coffee/coffee.html`), with the first ever webcam being used to determine when a fresh pot of coffee had been made [188].

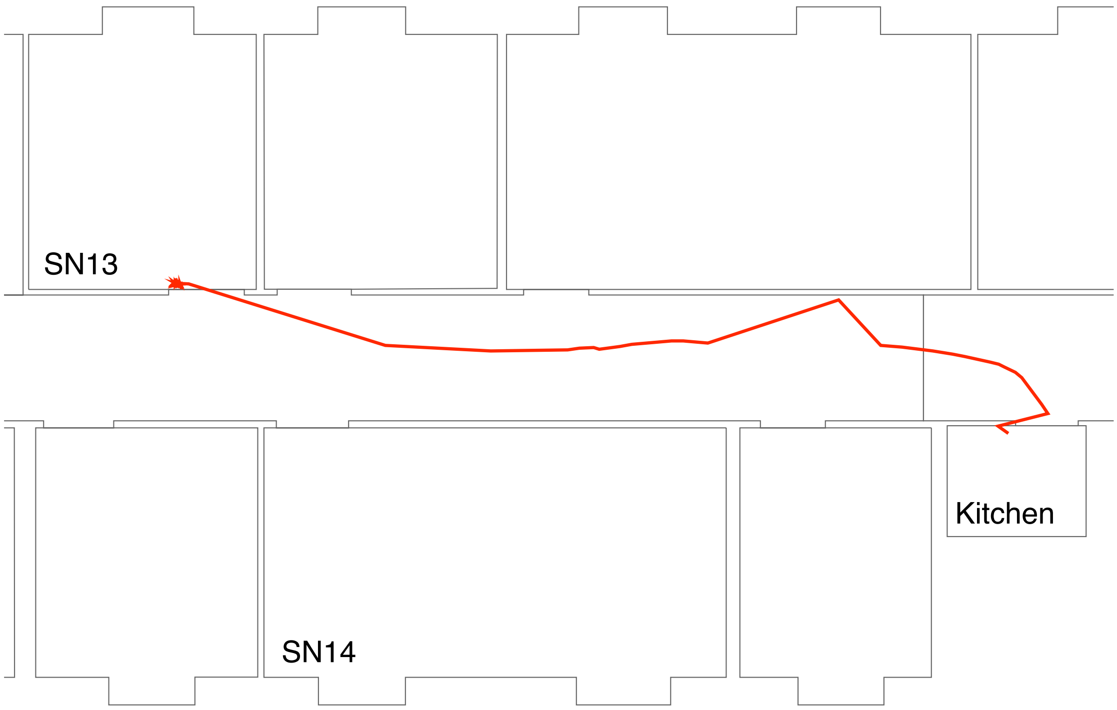

Figure 5.1: Location trace of walking to the coffee machine

for example, one straightforward option is to mine data from calendars [134] and for shared resources such as the Sentient Van or meeting rooms the online booking system provides the most reliable information on who is using it at any given time.

In general, however, one of the most valuable sources from which to infer context is location [43]. 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 5.1 shows the trace of a user walking from his office to the coffee machine recorded using the Bat system (see Section 2.7.2.1)—but note that since no sensors are installed in the kitchen, it is ambiguous exactly which device he was using. Matching locations against usage logs may reveal the user in most cases.

5.1.1 Low infrastructure location systems

Unfortunately, GPS positioning continues to struggle indoors due to the failure of satellite signals to penetrate buildings and is available for as little as 5% of a typical person's day [129]. To address this shortcoming, there have been many attempts at indoor location tracking, achieving a wide range of accuracies; these are surveyed in depth in Sections 2.7 and 2.8, and Table 5.1 summarises their key characteristics relevant to energy metering. While many systems can deliver impressive results the majority of these systems have not spread outside research labs, primarily due to the cost of deploying and maintaining

	custom base	custom target	cost	calibration	granularity
Infrared	✔	✔	medium	low	room
Ultrasound	✔	✔	high	high	cm
Radio	✔	✔	high	medium	cm
Inertial	✗	✔	high	low	cm
Sound	✗	✗	low	low	m–room
Power line	✔	✔	low	high	m–room
GSM	✗	✗	low	high	m
WiFi	✗	✗	low	high	m
FM	✗	✔	medium	high	m
DECT	✗	✔	medium	high	m
Bluetooth	✗	✗	low	low	m–room

Table 5.1: Summary of characteristics of location system technologies

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. Furthermore, all the additional equipment comes with an associated embodied energy cost.

A personal energy meter requires a reliable location technology that is viable for deployment across entire buildings using standard infrastructure and without extensive calibration. This rules out fine-grained tracking systems such as the Bat system (Section 2.7.2.1) since these require the retro-fit of dedicated infrastructure and the need to issue building occupants with custom tracking devices.[2]

Personal experiences with indoor location in many forms over many years have shown that even coarse location can provide useful location-aware applications. The Active Badge project was adopted so widely because the room-level location it provided was sufficient to enable a key application: finding co-workers quickly (see Section 2.7.1.1). This simple application demands only that users can be reliably located with room-level accuracy. The usage of that system waned because users forgot to carry their badges, didn't take responsibility for replacing batteries, and found that tracking failures were common (badges were easily obscured by clothing).

Precise location information is necessary to distinguish which of several nearby devices an individual is using, but even very coarse-grained data can be valuable for personal energy metering:

building-level location—knowing whether an individual is present or absent—is sufficient to enable the sort of apportionment described in Section 3.5. This is therefore the minimum required for energy metering, but even this provides a very significant improvement in the quality of data a personal energy meter can provide. Without it, only static apportionment, or apportionment based on standard working patterns, is possible.

[2]Experience shows that users tend to forget to wear their Bats. Additionally, there is a constant maintenance task in keeping Bats operational (especially with respect to battery power) and keeping the model of the world up-to-date.

area-level location—knowing which room, corridor or wing a person is in—additionally allows the energy consumed by shared resources, such as lighting, in that area to be divided among its occupants.

fine-grained location is necessary to determine the user of anonymous shared devices as described in the previous section.

A location system for the personal energy meter must also satisfy the following goals:

1. The system should be able to track users continuously
2. The system must be able to track multiple users simultaneously.
3. The system must be easily adopted by a very high percentage of building users and remain in regular use, avoiding the issue of users forgetting to enable tracking in some way.

The final criterion is one of the hardest to achieve. Mansley et al. explored a number of strategies for encouraging use of location systems, with particular reference to the Bat system, and discovered three main barriers to adoption [141]:

> A significant minority of people still do not regularly wear their Bat. An informal survey threw up a number of reasons: (1) apathy or forgetfulness—they just do not feel it is worth the effort or forget to put it on; (2) privacy—there are some who are concerned about their privacy, and so do not wear one on principle; and (3) discomfort—some people find wearing a Bat around their neck uncomfortable or annoying.

In order to get around the first and third of these problems, it makes sense to derive indoor location from the devices that people already carry. Mobile phones in particular are ubiquitous in many societies, feature an increasing array of communications technologies and are carried everywhere by their owners who find them useful and are motivated to ensure they remain charged and functional. Although it has been suggested that mobiles are sometimes farther from their owners than one might expect [162], they are more likely to be carried than any other device that could be used for location tracking.

To obtain in-building tracking using unmodified mobile telephones it is necessary to overlay location tracking on established technologies. Such tracking has been demonstrated using signals from the mobile telephony networks, using WiFi and using Bluetooth; these systems are surveyed in depth in Section 2.8. At present, positioning from the telephony networks is too coarse for in-building location (Section 2.8.4). WiFi-based tracking has received much attention in recent years (Section 2.8.5), but it is not ideal for general tracking. It is associated with high power demands (and the resultant difficulty in convincing users to leave handset WiFi turned on), difficulty in set up and maintenance, small market penetration of suitable (WiFi-enabled) handsets.

Bluetooth-based systems are particularly attractive as they have low power requirements by design and almost everyone already carries a mobile phone that supports Bluetooth and has a computer on his or her desk. This is therefore the option adopted here to build example location systems suitable for personal energy metering.

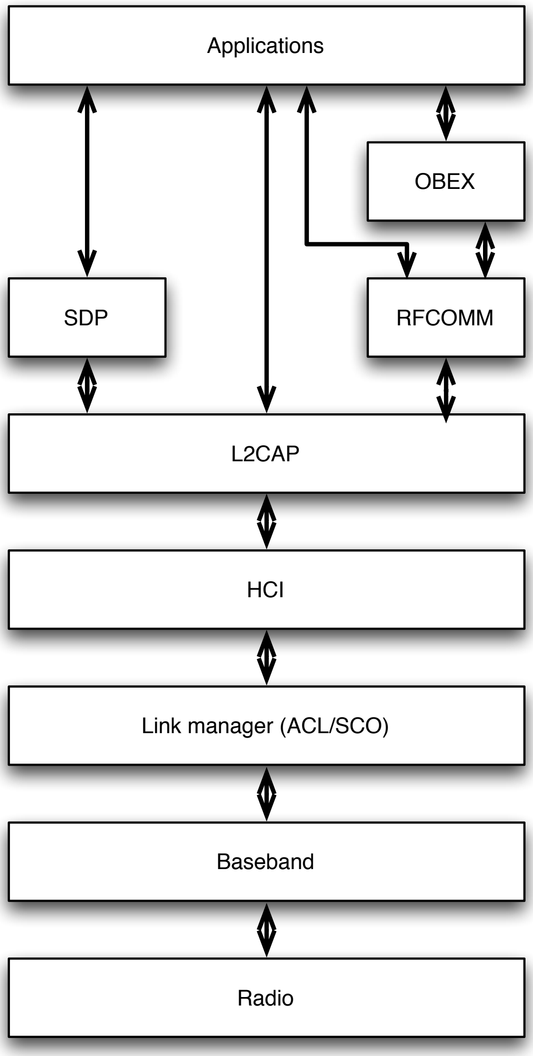

Figure 5.2: The Bluetooth protocol stack.

5.2 Bluetooth review

This section provides a brief overview of the parts of the Bluetooth specification relevant to the remainder of this chapter. Bluetooth was developed as a low-power wireless replacement for RS232 serial communication, and has become almost ubiquitous in mobile devices. The protocols are quite complex; the specification contains full technical information.[3]

5.2.1 Bluetooth connections

The Bluetooth standard defines three different types of connection that are layered to form the Bluetooth stack, shown in Figure 5.2. The most fundamental is the Asynchronous Connectionless Link (ACL). No more than one ACL can exist between any two Bluetooth devices at any given time, and it must be established before any other connection can be

[3] `http://www.bluetooth.org`

made. An ACL will disconnect only when no higher-layer connections have existed for a certain time (2 s seems to be the preferred value of current stacks).

Directly above the ACL is the Logical Link Control and Adaptation Protocol (L2CAP) layer. This is a packet-based layer that provides guaranteed packet sequencing and a selectable degree of delivery reliability. Once established, an L2CAP connection remains open until either end explicitly closes it, or the Link Supervision Time Out (LST) expires. The LST is the time for which communicating devices are out of range before the ACL connection is reported as destroyed. The default LST is 20 s in many stacks, but can be configured dynamically per-ACL.

Above L2CAP sits the Radio Frequency Communications (RFCOMM) layer, which is the reliable stream-based protocol used by most Bluetooth applications. It represents the type of connection most people mean by 'Bluetooth connection'.

5.2.2 Adherence to the specification

The Bluetooth specification is large, complex, and ever evolving. Implementations of it are not completely consistent in their interpretation and furthermore manufacturers have chosen to adapt some parts of the protocols to increase security. Only a few stack implementations are mature enough to be considered stable. This can lead to unpredictable behaviour, especially when using consumer devices with embedded Bluetooth stacks. A number of the mobile phones tested here exhibited occasional instability and unexplained Bluetooth behaviour, some of which appears to be intentional. The Nokia 6300, for example, appears to have a connection timeout associated with L2CAP connections—after 20 s, the ACL is dropped if no RFCOMM connection has been established.

5.2.3 Paging and inquiry

Bluetooth uses frequency hopping within the 2.4 GHz radio band for channel robustness. Every device retunes its transceiver every 625 μs to one of 79 Bluetooth channels. The precise sequence of changes is derived from its address and its local Bluetooth clock. For two devices to communicate they must have the same hopping sequence and phase at any given moment. To reach this state, Bluetooth defines a protocol that the two devices must follow. The protocol is known as *paging* and involves the master device sending search packets ('pages') addressed to the *slave* device until it replies. It is useful to consider the behaviour of the master and the slave separately; Figure 5.3 illustrates the process.

Slave. The slave device periodically listens for page requests on a particular radio channel for 11.25 ms. A total of 32 channels are used for paging, and the frequency the slave listens on is changed every 1.28 s according to a sequence also derived from its clock and its device address. If the slave does not detect a page, it sleeps for set period of time, T_{pg_scan}. The Bluetooth specification defines three *SR* modes which a device can adopt and which provide a limit on T_{pg_scan}. These modes are R0 ($T_{pg\_scan} = 0$), R1 ($T_{pg\_scan} \leq 1.28$ s) and R2 ($T_{pg\_scan} \leq 2.56$ s). The default mode is R1, but some mobile devices adopt R2 to save power.

Master. With each page sent out, the master needs to estimate the frequency that the slave will be listening on. The device's address tells it the hopping sequence, but it

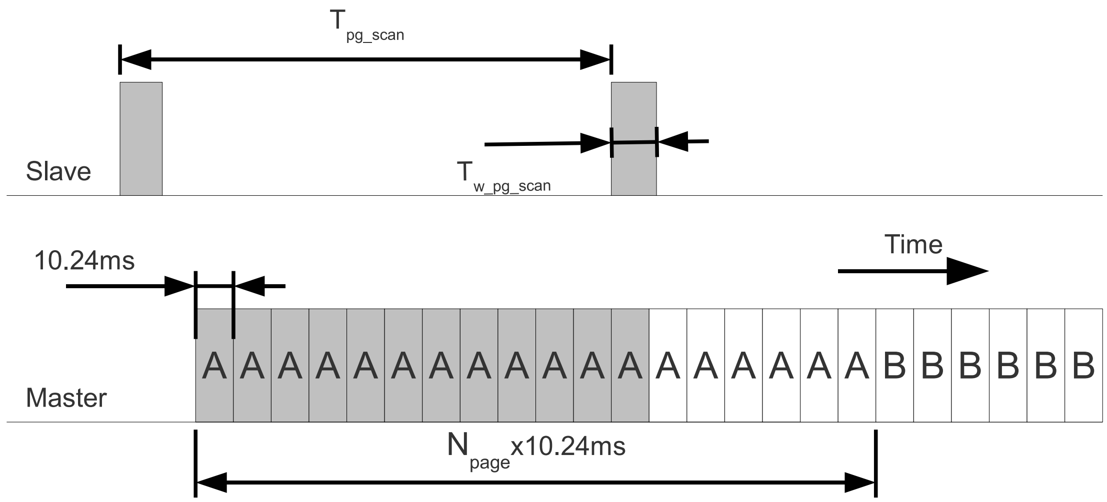

Figure 5.3: The paging process with parameters marked. A slave periodically listens on a single frequency for $T_{w\_pg\_scan}$. When paging, the master pages the 16 A frequencies in turn, each cycle taking 10.24 ms. After N_{page} cycles, it repeats using the B frequencies. In this example, shaded cycles for the master indicate the cycles needed to connect to the slave shown [87].

can only know the current sequence position by having an estimate of the Bluetooth clock on the slave. Once it has such an estimate, it then computes the most likely frequency to transmit on. However, in the 11.25 ms that the handset listens for, it has time to send pages to 16 channels. Thus it chooses a set of 16 frequencies that are adjacent in the hopping sequence and centred on its best estimate of the listening frequency. This set of 16 frequencies is known as 'train A' and allows for a degree of error in the estimate of the slave's clock. It then cycles over this train of frequencies continuously for some for $T_{window} \geq T_{pg\_scan}$ seconds. Assuming the slave is listening on a train A frequency, it will hear the page and respond, implicitly syncing the two devices. On average this should take $\frac{1}{2}T_{pg\_scan}$ to complete.

5.2.4 Radio interference

Whenever a device is paging or inquiring, it is flooding the 32 channels used for those purposes. In fact, the specification recommends (but does not require) that any established connections be temporarily parked during a page or inquiry. Inevitably established connections will be disrupted by continuous scanning or paging, meaning that tracking in this manner may limit the use of Bluetooth as the *communications* medium it was intended to be.

This was evaluated using three machines: one a nominal master; one a slave; and one used to provide interference in the same area (a nominal 'external' machine). Figure 5.4 depicts how long it took to transfer a 1 MB file from the master to the slave and back again under different conditions, which showed that external scanning or paging had little effect. However, when either the master or slave were continuously scanning or paging, the data throughput for the established connection fell dramatically. The connection was

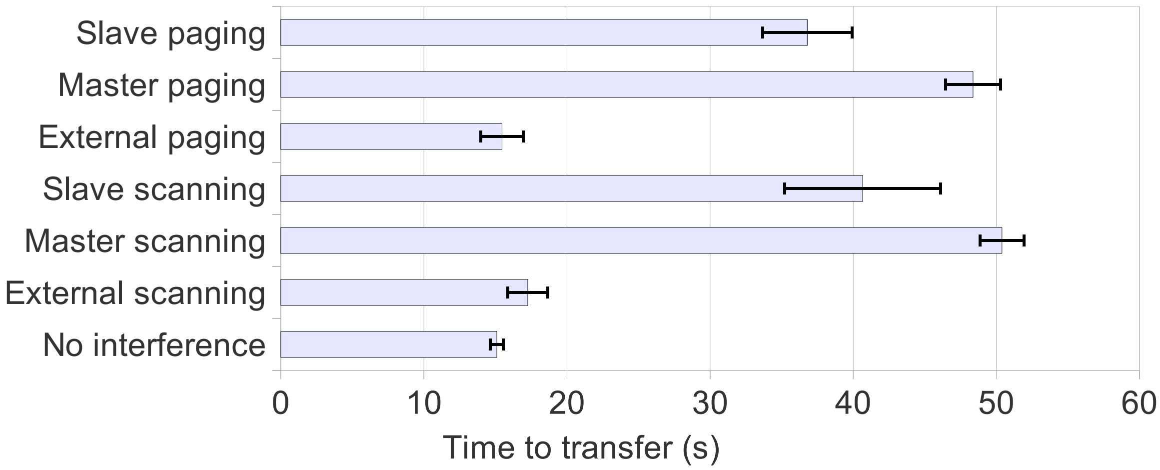

Figure 5.4: Measuring connection disruption. The chart shows the mean time taken to transfer a 1 MB to and from a Bluetooth device whilst the master, slave or an external device was continuously scanning or paging a fictional device address. The error bars show ±1 standard deviation.

slowed rather than severed or permanently parked. These data apply only to the BlueZ stack, but indicate that the tracking techniques may be able to co-exist with 'normal' Bluetooth usage, albeit reducing the data throughput.

5.3 Scan-based tracking

As discussed in Section 2.8.8, most existing Bluetooth-based tracking systems locate a device using the inquiry, or scan, mode of Bluetooth. In this mode, a base station transmits a discovery packet on each of 32 radio channels. Devices set to 'discoverable' respond to this packet, identifying themselves. However, the response follows a random delay in order to minimise the chance of response collisions when multiple devices respond. The result of this protocol is that an inquiry must run for 10.24 s to detect reliably all devices in range (and longer still if the radio conditions are unfavourable).

Such a system has two possible configurations. Initially, many systems tracked mobile devices by continually issuing inquiry packets from a network of fixed beacons (hereafter referred to as Base scans Target, or BsT). This has the advantage that no custom code need be deployed on the handset, but is often perceived as a privacy risk since anyone can track a handset by creating their own network of beacons. Additionally, handset manufacturers are increasing security by changing the 'discoverable' mode to be a time-limited handset state; this is the case for the iPhone[4] and Android.[5] The system must therefore regularly renew the discoverability request on the target. This introduces at least two problems: firstly, the device may miss an inquiry packet in the gap between discoverability being disabled and re-enabled, and secondly, it must be running custom software to perform the renewal which limits universality and complicates deployment.

[4]up to and including at least iOS 4.2
[5]up to and including at least Android 2.2

More recent tracking attempts have therefore concentrated on the mobile handset scanning for the fixed beacons (Target scans Base, or TsB) [27, 79, 100]. This is more secure since Bluetooth does not require that a scan packet identify its source address. However, it requires custom application code on the handset, will typically draw more power, and requires a data communications channel to publish positions. Regardless of these issues, both schemes have traditionally suffered from further problems:

High tracking latency. Since each scan can take 10.24 s, these systems have a very slow update rate that does not support dynamic tracking.

Devices in the system must be discoverable. In order for a scan to locate a nearby device, that device must be discoverable. It therefore announces itself to the world, and becomes a target for hackers regardless of whether it is a handset or a beacon.

This chapter investigates potential improvements that could create a Bluetooth-based tracking system suitable for personal energy metering.

5.3.1 Update rates

Unlike a page, an inquiry does not complete once a response has been heard. Instead, most Bluetooth systems require that an inquiry request specifies the duration of the inquiry. Some stacks support the option of terminating a scan once a specified device is observed, but in tests this has merely prevented subsequent events from being passed up to higher levels; a new scan cannot be started until the hardware has completed the last one and the duration of that must be set at the start of the inquiry. Therefore the update rate for a scan-based system would seem to be determined by the chosen scan duration.

Figure 5.5 shows the effect of varying T_{inq_scan} on the time taken to spot a particular device. The data were collected by continuously scanning for five minutes (30 10.24 s scans) with each T_{inq_scan} value and using the BlueZ HCI event API to record when responses were received. The device was successfully found in all 30 scans at each setting, and always within the first 5.12 s. The CDF shows that even when small T_{inq_scan} values are in use, the device may be missed in the first inquiry train, and thus each scan must last at least 3 s regardless of the T_{inq_scan} value. This bounds the update rate to 0.333 Hz.

However, since a listening slave cannot identify the source of an inquiry by design, it must respond to *every* inquiry it receives. It is possible that, having already responded to an inquiry, a slave continues to listen and hears and responds to another of the inquiry requests *from the same inquirer*. Thus each scan can result in multiple responses from the same device. Many stacks allow applications to register with the HCI layer for all response events, even repeats. This gives an effective update rate that is higher than the inquiry rate.

Care must be taken in using this property. Not hearing from a particular device during an inquiry is strong evidence that it is out of range. However, not hearing from a device multiple times within a single inquiry is *not* strong evidence that it has moved out of range during the inquiry; merely that the transmitter and receiver did not sync their frequencies again. Nonetheless, the repeated responses are particularly useful for a fingerprinting system, since each response may be set to carry an independent RSSI estimate. Then

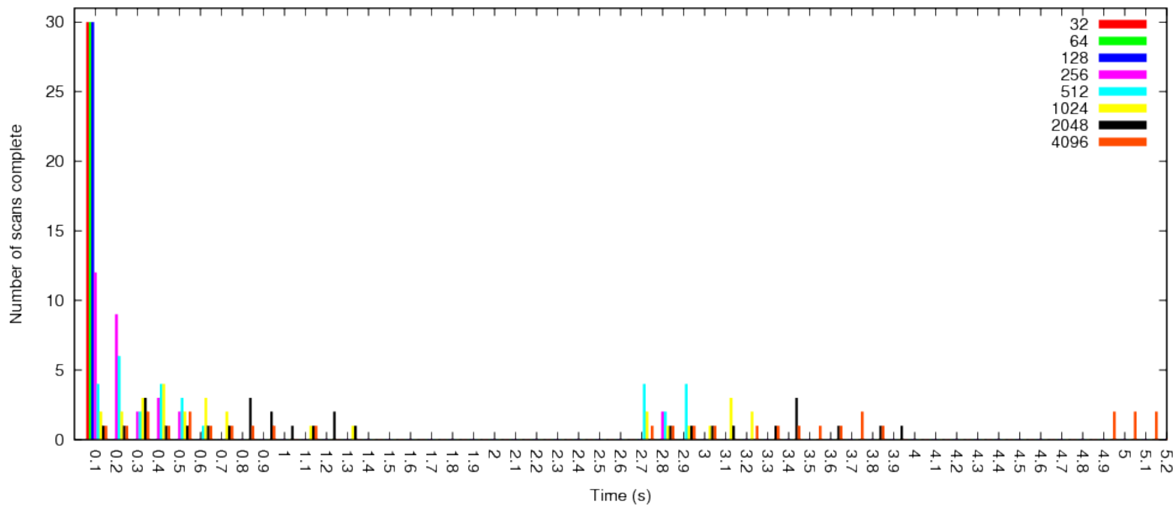

(a) Histogram of times

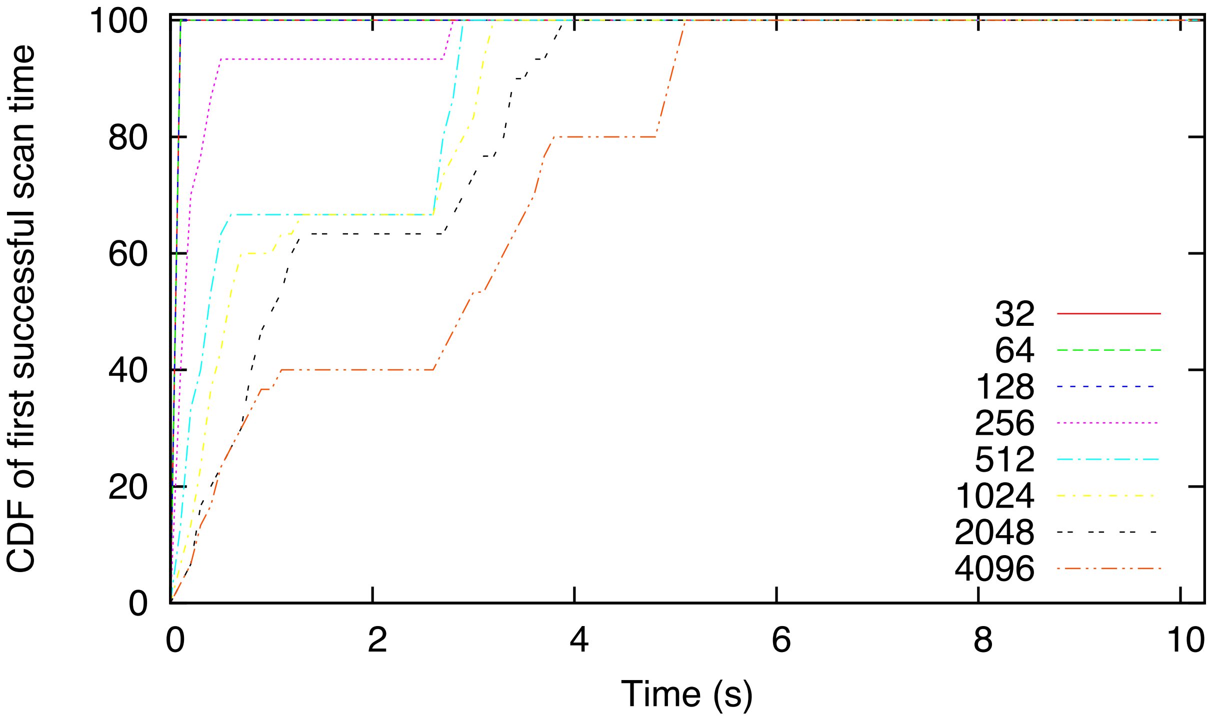

(b) CDF of times

Figure 5.5: Analysis of the inquiry time before first finding discoverable device whilst varying T_{inq_scan} (shown in number of slots).

a single inquiry can produce multiple RSSI estimates to multiple devices, allowing for better matching.

Figure 5.6 illustrates the distribution of all HCI inquiry events associated with the data used for Figure 5.5. There were nearly 2,000 scan events (each with an independent RSSI estimate) generated across the 30 scans when T_{inq_scan} was set to 32. This implies an average update rate of almost 7 Hz although these results were tightly clustered at the start of each 2.56 s interval. This is merely a reflection that there is a higher chance of the master and slave matching frequencies shortly after they have just done so before.

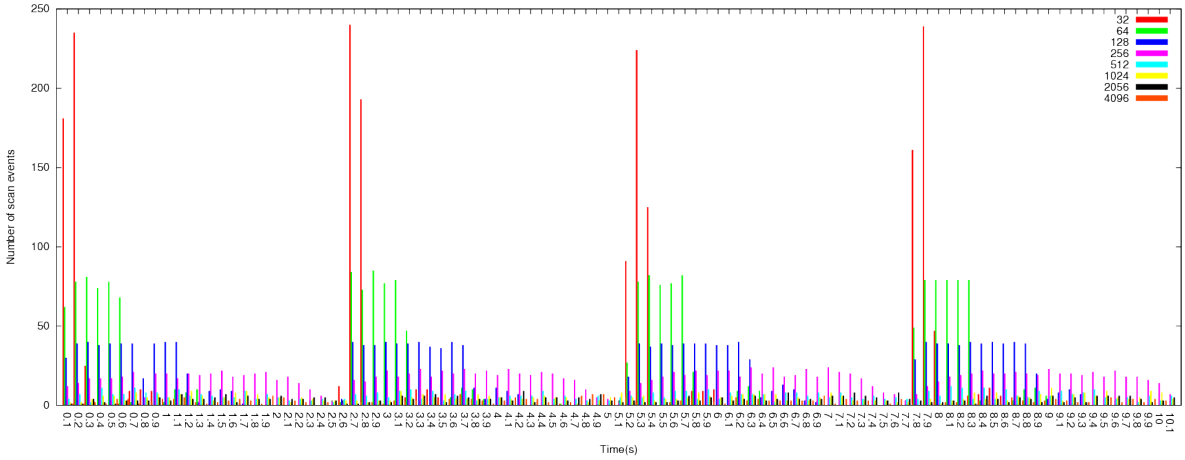

(a)

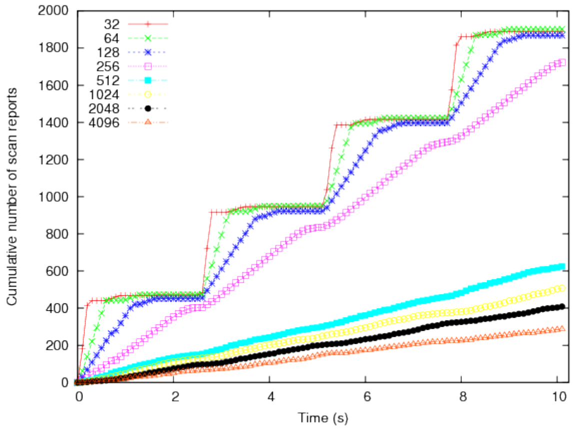

(b)

Figure 5.6: Analysis of all inquiry events corresponding to one device whilst varying T_{inq_scan} ($T_{w\_inq\_scan}$ =18). (a) Histogram of sighting times (b) cumulative number of sightings made.

However, for $T_{inq\_scan} \geq 512$ there is a more uniform distribution of events, with a regular update rate of up to approximately 1.7 Hz. Even the default value of $T_{inq\_scan} = 4,096$ can provide an regular update rate of 0.95 Hz using this approach. These results are very useful for fingerprinting.

Figure 5.7 shows also the effects of varying $T_{w\_inq\_scan}$ whilst holding T_{inq_scan} at its default of 4,096 slots. As expected, there is a general trend of faster/more inquiry events with a larger window, although the random nature of the inquiry procedure means that nearby values of $T_{w\_inq\_scan}$ may not behave as expected in some circumstances. In Figure 5.7, for example, $T_{w\_inq\_scan} = 32$ gave an unexpectedly high update rate; this can be attributed to fortuitous alignment between the hopping sequences. Nonetheless, the general trend is evident and, if consuming all inquiry events, average update rates of between 0.4 Hz and 1.1 Hz were observed.

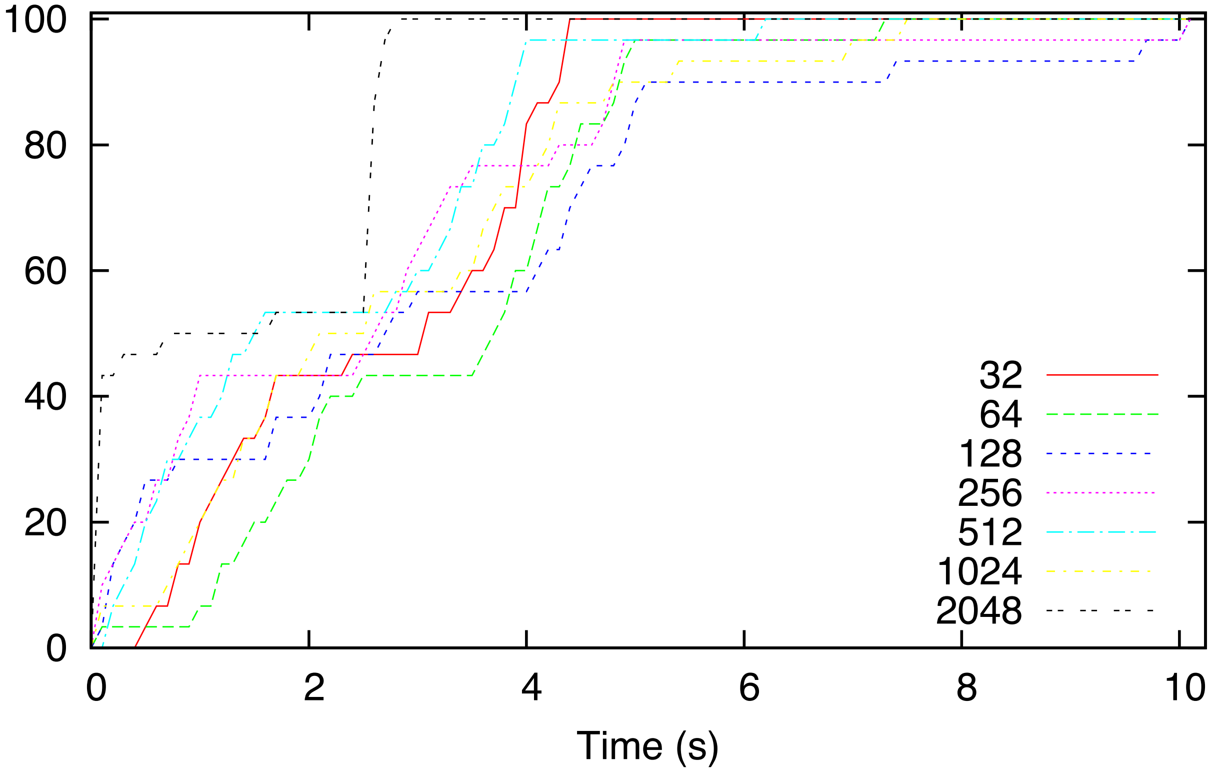

(a)

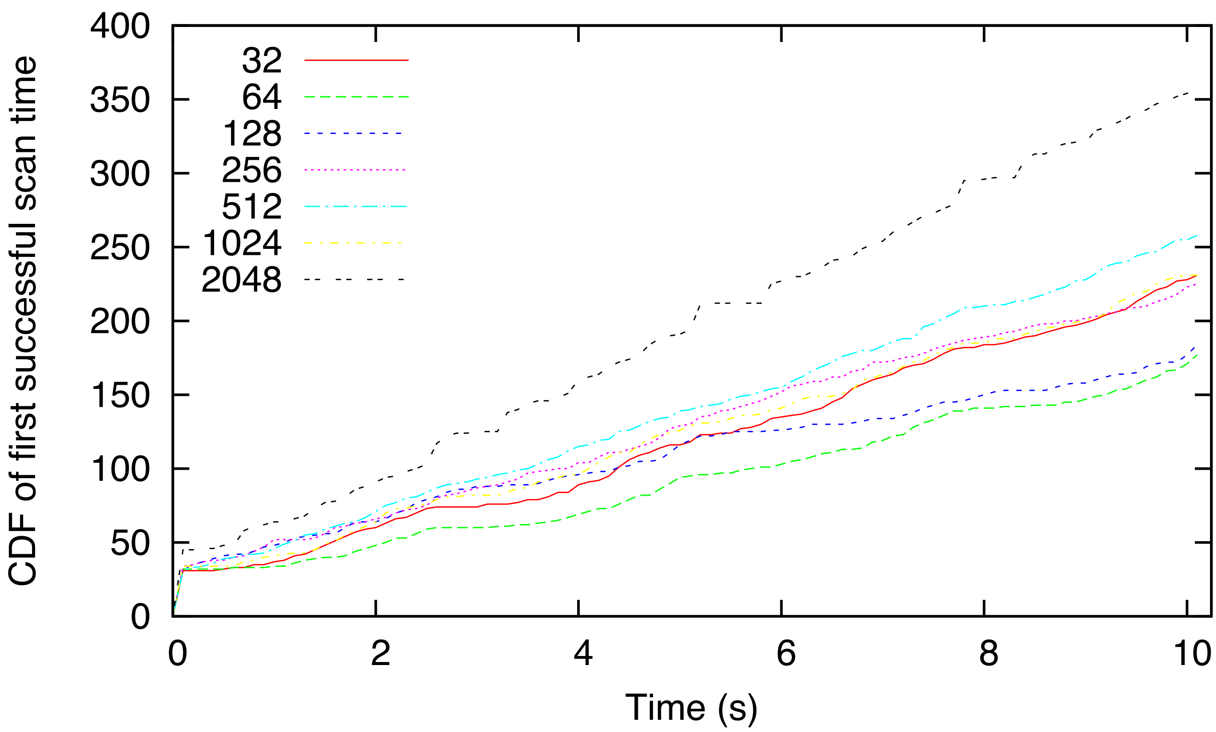

(b)

Figure 5.7: Analysis of inquiry events whilst varying $T_{w\_inq\_scan}$ ($T_{inq\_scan} = 4,096$). (a) First inquiry events only (b) All inquiry events.

5.4 Connection-based tracking

Scan-based tracking detects proximity between handsets and fixed beacons by continually scanning for all devices. An alternative, novel notion described here is *connection-based* tracking, whereby two specific devices are characterised as proximate if one can connect to the other i.e. if beacon B can connect to handset H then B and H are proximate. Either the target initiates the page (Target connects Base, TcB) or vice-versa (Base connects Target, BcT). It is possible to track a moving target by continuously attempting to connect and immediately disconnect in a manner analogous to the standard ping tool. If the target is static, however, it may be preferable to maintain a single connection, monitoring its properties to determine when the target is moving again.

In many senses this is notionally similar to an scan-based system where each inquiry targets a specific device rather than every device. This has the obvious disadvantage that multiple such inquiries would be needed to find all the local devices. However, the specificity also allows devices to be targeted quickly rather than with extended broadcasts i.e. a single inquiry should complete faster, allowing for multiple inquiries.

The key motivation for investigating connection-based tracking is that its Bluetooth implementation can potentially address the scan-based shortcomings identified. The remainder of this section is devoted to characterising both theoretically and experimentally the relevant properties of Bluetooth that permit this connection-based tracking.

5.4.1 Connection authorisation

Connection-based systems may be based on one of several different connection types. At a fundamental level all of them use an ACL (see Figure 5.2), but implementation constraints mean that the choice of connection type may be important.

5.4.1.1 RFCOMM

The biggest obstacle to connection-based tracking using RFCOMM connections is the requirement for explicit *pairing* of every handset with every fixed base. The pairing process is governed by the Bluetooth security manager and involves creating a shared secret (a passphrase or PIN) between the two devices. Because each pairing operation requires human input, it is not generally practical to pair every handset with a large deployment of bases. However, one novel workaround is possible. Although platform-specific, easy pairing *is* achievable if the base addresses are spoofed to the same address and the BlueZ stack is used as follows:

Handset	Description
T-Mobile G1	Android 1.1
Apple iPhone 3G	iPhone OS 2.2
Nokia 6300	Series 40 3rd Edition, Feature Pack 2
Nokia N80	Series 60 3rd Edition (Symbian OS 9.1)

Table 5.2: The test handsets used

1. Pair the target with a particular base.

2. Copy the relevant entry in the linkkeys file (found in /var/lib/bluetooth/) on the base to the equivalent file on every other base.

3. Using address spoofing, set every base to have the same address as the base used in pairing.

Note that a failure to pair (rather than a failure to connect) may be sufficient to indicate presence, avoiding any need for the pairing attempt to be successful. Unfortunately, few stack APIs support a distinction between failure to pair and failure to communicate. Furthermore most mobile devices will automatically initiate a graphical pairing prompt under these circumstances rather than fail silently.

If addresses are spoofed, an alternative method must be used to distinguish between bases such as hostnames transmitted over the connection or individual logs kept by each base. This is discussed further in Section 5.4.6.

5.4.1.2 L2CAP

There are some communications services that do not usually require this level of authentication. In particular, the creation of an ACL and a basic L2CAP connection is *almost* universally authorisation-free. Although the resultant connections are limited in use for communications (they support little more than low-level testing) they are sufficient for tracking usage because if they are successfully established, the two devices must be within range of each other. Additionally, the low-level tasks they do support, such as RSSI measurement and L2CAP echo requests (directly analogous to the familiar ICMP ping packet in IP), allow continual monitoring of the connection. However, not all Bluetooth stacks feature explicit L2CAP support—this is a particular issue for embedded devices such as mobile phones, meaning that while L2CAP is a viable option for BcT systems it is less practical for TcB systems.

One possible workaround is the use of the Service Discovery Protocol (SDP) channel. Most general purpose Bluetooth devices use an SDP server to advertise the logical channels linked to the services they offer. SDP itself is built on an L2CAP connection with a preset channel number of 1. Therefore it may be possible to issue an SDP lookup for a non-existent service to a device. If the stack API in use differentiates between failure to connect to the SDP server and failure to lookup the service, it is possible to infer locality of the device or otherwise.

5.4.1.3 LM name request

As discussed above, any L2CAP or RFCOMM connection implies an underlying ACL connection. However, the lifetime of the ACL is not matched to the lifetime of the logical channels that prompted its creation. Since the cost of a page is high, the ACL is maintained for a short period of time in case the application wishes to re-establish a logical connection to the remote device. This time is known as the *ACL timeout*. In the context of tracking, care must be taken to ensure that any requested disconnect at the logical level results in a disconnect at the ACL level before proceeding to reconnect to the same address.

Unfortunately, few stacks expose the ACL to the developer. Instead, the *remote name request* feature in the Bluetooth specification can be used. This function queries the human-readable name of a remote device. To do so it first establishes a connection (via standard paging), exchanges the necessary data and then, crucially, disconnects *immediately*, so maintaining the ACL for only a minimal length of time.

BlueZ provides the `HCI_read_remote_name()` function to initiate such a name request from the HCI layer. This has been used in the RedFang system to detect the presence of Bluetooth devices that are not discoverable, connected or paired to the enquirer but has not previously been used for tracking.[6]

The initial motivation for using low-level L2CAP connections was to avoid the Bluetooth pairing procedure, which is often invoked automatically for higher-level connection protocols such as RFCOMM. Unfortunately, some manufacturers are now creating consumer devices that use Bluetooth security mode 3. In this mode, pairing is required for the establishment of a full ACL, and therefore needed for L2CAP connections too.

However, the temporary ACL established through the `HCI_read_remote_name()` call may *not* be subject to authentication for a device in security mode 3. This appears to be true for the BlueZ stack in particular. The use of the name request therefore remains a viable option for tracking.

Alternatively, the distributed pairing technique described in Section 5.4.1.1 can be used to allow L2CAP connections to be used for tracking when pairing is required. However, there is then little advantage of L2CAP over RFCOMM—indeed, the latter is generally better supported.

5.4.2 Connection time

A connection-based tracking system will constantly attempt connections between devices. The expected latencies in tracking will be dictated by the connection times, which are quantified here.

If the master has a good estimate of the slave's clock, it should be able to connect in train A (see Section 5.2.3). On average this should take $\frac{1}{2}T_{pg\_scan}$ to complete.

The master normally obtains the important clock estimate by completing a lengthy inquiry process. If it does not have an estimate (e.g. inquiry is avoided as in the connection-based technique) it is forced to centre train A on a randomly chosen paging frequency. There is now a 50% chance that the slave will not be listening on a train A frequency. After T_{window} seconds without a reply, the master repeats the entire process using the other 16 frequencies; 'train B'. On average, the pager will have to use train B half of the time. Thus, the average connection time (ignoring train changes) is

$$\begin{aligned} T_{conn} &\approx \frac{1}{2}.\frac{1}{2}T_{pg\_scan} + \frac{1}{2}(T_{window} + \frac{1}{2}T_{pg\_scan}) && (5.1) \\ &\approx \frac{1}{2}(T_{pg\_scan} + T_{window}) && (5.2) \end{aligned}$$

[6] `http://www.securiteam.com/tools/6L00K008LE.html`

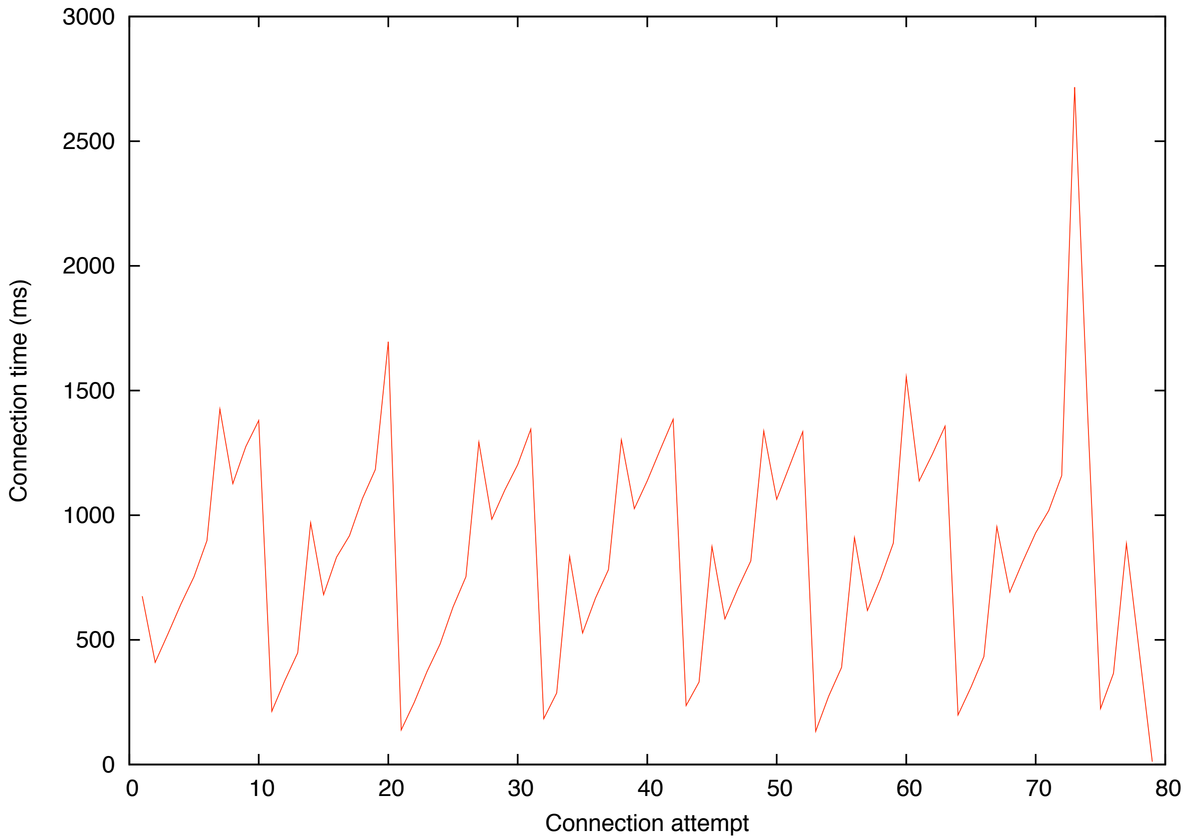

Figure 5.8: Experimentally measured L2CAP connection times for the iPhone [87]

and $T_{pg\_scan} \leq 2T_{window}$.

Therefore for default values of $T_{pg\_scan} = 1.28$ s and $T_{window} = 2.56$ s a successful page will take 0.64 s on average in train A, and 1.92 s on average otherwise. In the worst possible case, a page should complete within 5.12 s, which is still twice as fast as a default inquiry.

Once a connection is established, the master can cache the clock offset for the slave and use it for subsequent connections to the slave. Over time, however, the clocks drift apart and the estimate will degrade.

None of this analysis, however, incorporates any setup overheads at the master or slave once paging has succeeded. One would expect there to be a systematic increase to the connection times. Figure 5.8 shows the connection times to an Apple iPhone. For each datum, a new connection was created and its creation time measured using the UNIX C function `gettimeofday()` on the connecting system. The start of each connection was spaced by 5,000 ms, chosen to be indivisible by 1,280 ms.

Since the iPhone has $T_{pg\_scan} = 1.28$ s, each subsequent connection attempt would start at a different point in the page scan cycle of the iPhone.[7] Ignoring connection overheads, the connection time should move linearly between zero (connection starts just as the slave's page cycle does) and 1,280 ms (the connection starts just after the cycle ends). The experimental connection times incorporate an additional offset of approximately 200 ms, which is the connection setup overhead. Additionally, one connection attempt failed altogether, resulting in a connection time of two page slots.

[7]Note that BlueZ cached the clock offset so all connections are expected to occur in train A

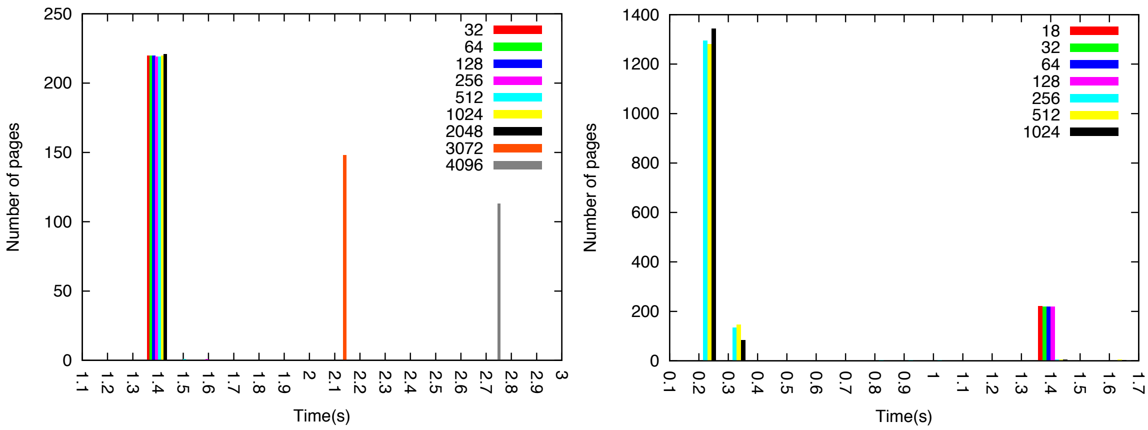

(a) (b)

Figure 5.9: Distribution of paging times for (a) $T_{w\_pg\_scan} = 18$, varying T_{pg_scan}; (b)varying $T_{w\_pg\_scan}$, $T_{pg\_scan} = 2,048$.

Ideally, connection would be near instantaneous for connection-based tracking. Figure 5.8 shows that with unmodified handsets, it is realistic to expect connection times to be bounded by about 1.5 s. Until a connection is established or 1.5 s has elapsed, a given beacon will not be able to verify whether or not a handset is in range. The equivalent figure for scan-based searching is 10.24 s, although as noted previously, multiple handsets can be queried simultaneously.

This time can be reduced if the paging parameters can be tuned. This is not often the case on mobile devices without modification, but is generally possible on PCs. As can be seen from Figure 5.3, there are three parameters that can be used to minimise the duration of a paging attempt:

- reduce T_{pg_scan} on the slave;
- increase $T_{w\_pg\_scan}$ on the slave;
- reduce N_{page} on the master.

Potential update rates were evaluated using two machines, each with a Belkin Mini Bluetooth adaptor, separated by a metre or so of free space (i.e. no interference expected). Both machines ran a modern Linux kernel[8] and used the BlueZ stack. BlueZ allows dynamic tuning of the T_{pg_scan} and $T_{w\_pg\_scan}$ parameters but not N_{page}. Similarly, the analogous inquiry parameters can be set, with $N_{inquiry}$ hard coded.[9]

Figure 5.9 shows the result of varying T_{pg_scan} and $T_{w\_pg\_scan}$ on the duration of an LM name request, which is dominated by the paging time. The durations were remarkably consistent for each setting. Update rates of almost 5 Hz were achievable by increasing the $T_{w\_pg\_scan}$ value; more conservative settings gave approximately 0.7 Hz; and the worst case update rate observed was approximately 0.35 Hz.

[8]Linux 2.6.32

[9]Technically the Bluetooth specification allows for no other value for $N_{inquiry}$ than 256. However lower values do not break functionality.

5.4.3 Disconnection time

Bluetooth connections can be either be shutdown in an orderly fashion by either end or severed when out of range. For the latter, there are two separate relevant timeouts, introduced in Section 5.2.1:

ACL timeout. This timeout is the amount of time that an ACL waits before disconnecting itself after all higher level logical connections are closed, as discussed in Section 5.4.1.3 above. The default timeout is 2 s. Some stacks provide tools to forcefully disconnect an ACL (BlueZ being one), but these are not standard.

Link Supervision Timeout (LST). When two Bluetooth devices have their communication disrupted, the Link Supervision Timeout (LST) defines how long they will wait for the other party to return before a disconnection event is issued to higher layers. Bluetooth sets the default LST at 20 s, although it can be changed dynamically once a connection is established. The LST is per-link, not per-device, so setting it at one end of the connection sets it for the other.

The issue with the ACL timeout has already been mentioned: if a host simply disconnects and immediately attempts a reconnect, the ACL timeout will not have been triggered and the baseband connection will be maintained rather than re-established.

In the context of tracking, the LST is triggered whenever the device goes out of range without first negotiating a disconnection. A complete disconnection then occurs only after the link supervision timeout (LST) has elapsed without communication. However, experimenting with test phones showed that when a connection was lost due to range, the connection did *not* re-establish itself if the handset re-entered the radio range before the LST completed. This means that if a handset leaves communications range of a host and returns within the LST, the original host may not be able to 'see' it until the LST has expired.

The connection-based techniques described continuously connect and disconnect, and minimising the time between these two events reduces the chance of triggering the LST since there is less opportunity for a device to move out of range. The remote name request is the optimal choice here, since it maintains an ACL for as little time as possible.

Another approach is to reduce the LST once a connection is established. If this is done immediately after a successful page, there is a relatively small window within which the default LST can be triggered. This window does exist, however, so an LST trigger is always possible. Note also that a reduced LST value will negatively impact the reliability of any other data connections between the two devices.

In practice the use of the remote name request is sufficient to avoid triggering the LST anything but irregularly.

5.4.4 Connection monitoring

Once a handset and beacon are connected, a forced disconnection signals that they are no longer co-located. Whilst this is useful information, it may be possible to infer more by monitoring the connection quality.

Handset	Echo process rate (Hz)
T-Mobile G1	13.00
Apple iPhone 3G	20.14
Nokia 6300	21.97
Nokia N80	20.26

Table 5.3: Experimentally measured ping rates for different handsets.

In general there are three metrics used to monitor a Bluetooth connection: Received Signal Strength Indicator (RSSI), Link Quality (LQ) and echo response time. The standard does not require a Bluetooth chip to report the RSSI or LQ, although most do. It is safe to assume that the Bluetooth chip of the beacons, where the metrics are to be collected, can be chosen to support these features.

The advantages of using the RSSI or LQ measurements are twofold. Firstly, they do not involve sending extra data wirelessly and so do not consume power at the mobile handset. Secondly, experiences indicate that many manufacturers update the reported RSSI value at 1 Hz or even faster, allowing for fast monitoring of a link if so desired.

An alternative method to monitor the connection is to use the round-trip time for an echo packet, similar to an ICMP ping in TCP/IP networks. Poor connections result in packet loss and retransmission and corresponding increases in round-trip time. Table 5.3 shows the experimentally-measured maximum rate of echo requests for each of the test handsets when close to the sending host. All testing was carried out using a desktop Linux machine (running Ubuntu 8.04 with the BlueZ 3.26 Bluetooth stack) with an attached class 2 generic Bluetooth 2.0 USB dongle and a range of mobile telephones taken as representative of those available on the market (see Table 5.2).

The round-trip time is of the order of 40 ms for a strong connection, so update rates faster than 1 Hz are easily achievable. However, this approach does require the handset to be more active and therefore reduces its battery lifetime; this is described further in Section 5.7.

5.4.4.1 RSSI/LQ and proximity

Although these metrics can be read, it is not immediately clear how useful they are. There is a growing body of work on the use of RSSI values for location determination through fingerprinting techniques (see Section 2.8.3.3). The ideas described so far could be used to gather RSSI readings at a higher rate than the usual approach (using inquiries) and so may improve those methods. However, fingerprinting requires detailed surveying (and continual re-surveying) which does not meet the goals of minimal maintenance and so it is not pursued here.

There have also been attempts to relate the RSSI or LQ values to absolute distances with which to derive position estimates [76]. Such approaches have demonstrated only limited success to date; much coarser use of the RSSI seems prudent.

A series of offices were surveyed using an iPhone to measure RSSI and a co-located Bat to measure locations simultaneously to within a few centimetres. The office wing was representative of a typical situation where Bluetooth had to co-exist with WiFi, walls attenuated signals and people unrelated to the experiment were permitted to work as

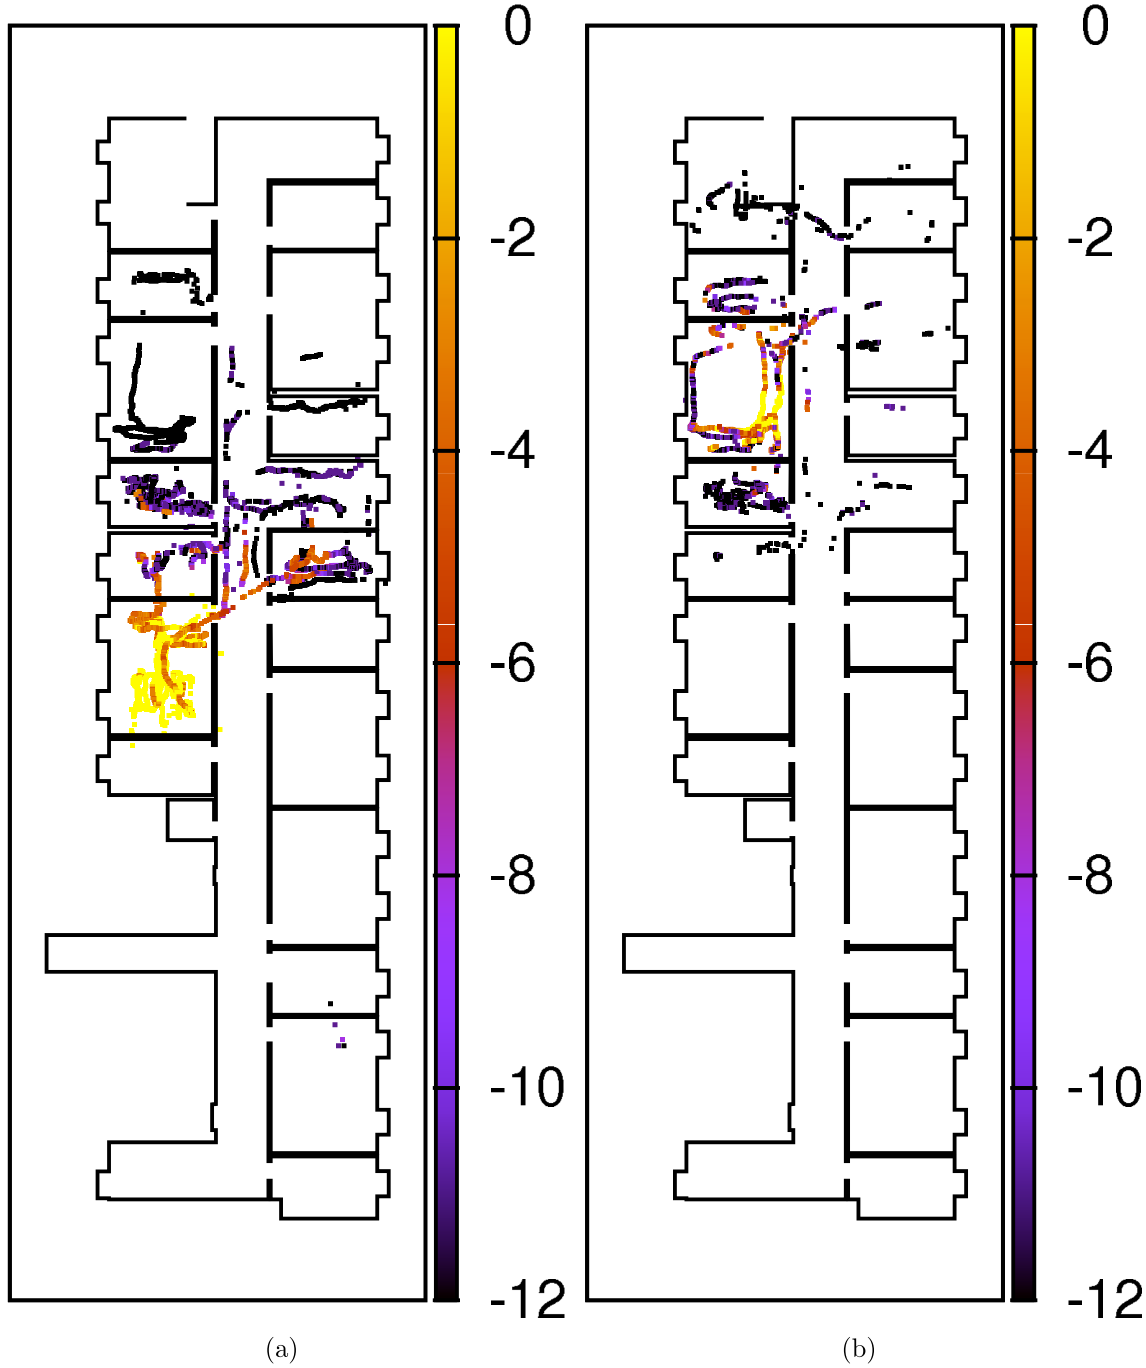

Figure 5.10: Experimentally measured RSSI values from four hosts, plotted against locations recorded using the Bat system [87]

usual. The traces from four hosts are shown in Figure 5.10. Figure 5.11 shows the observed relationship between the reported RSSI and the distance from the fixed Bluetooth master, based on data collected for 10 different masters. Note that an RSSI of -13 is actually a disconnection. Madhavapeddy and Tse performed further studies of Bluetooth propagation in the same offices using a similar technique [137] (see Section 2.8.8). It is clear from Figure 5.11 that there is no deterministic relationship between distance and RSSI, but that there is a qualitative trend (larger distances are associated with more negative RSSI values). This trend is used simply to associate a falling RSSI reading with

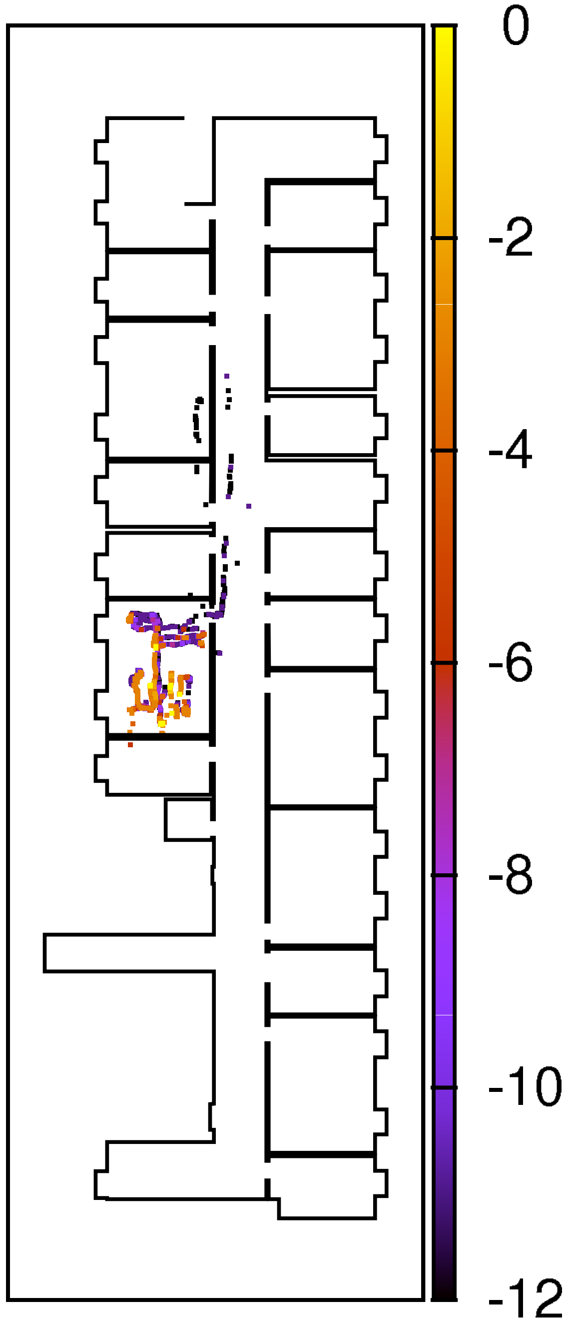

(a)

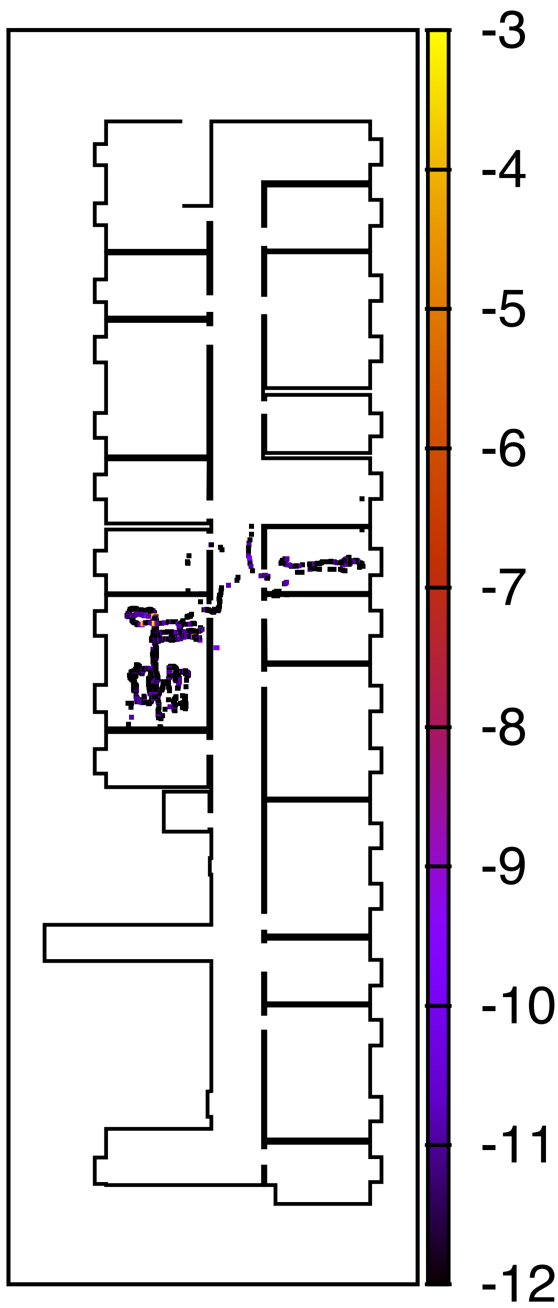

(b)

Figure 5.10: continued

a handset leaving the area, which is generally sufficient for these purposes.

5.4.5 A base-connects-target tracking system

Having characterised the fundamental properties of Bluetooth in the context of connection-based tracking, the next issue that arises is how to exploit these properties to build a wide-area tracking system.

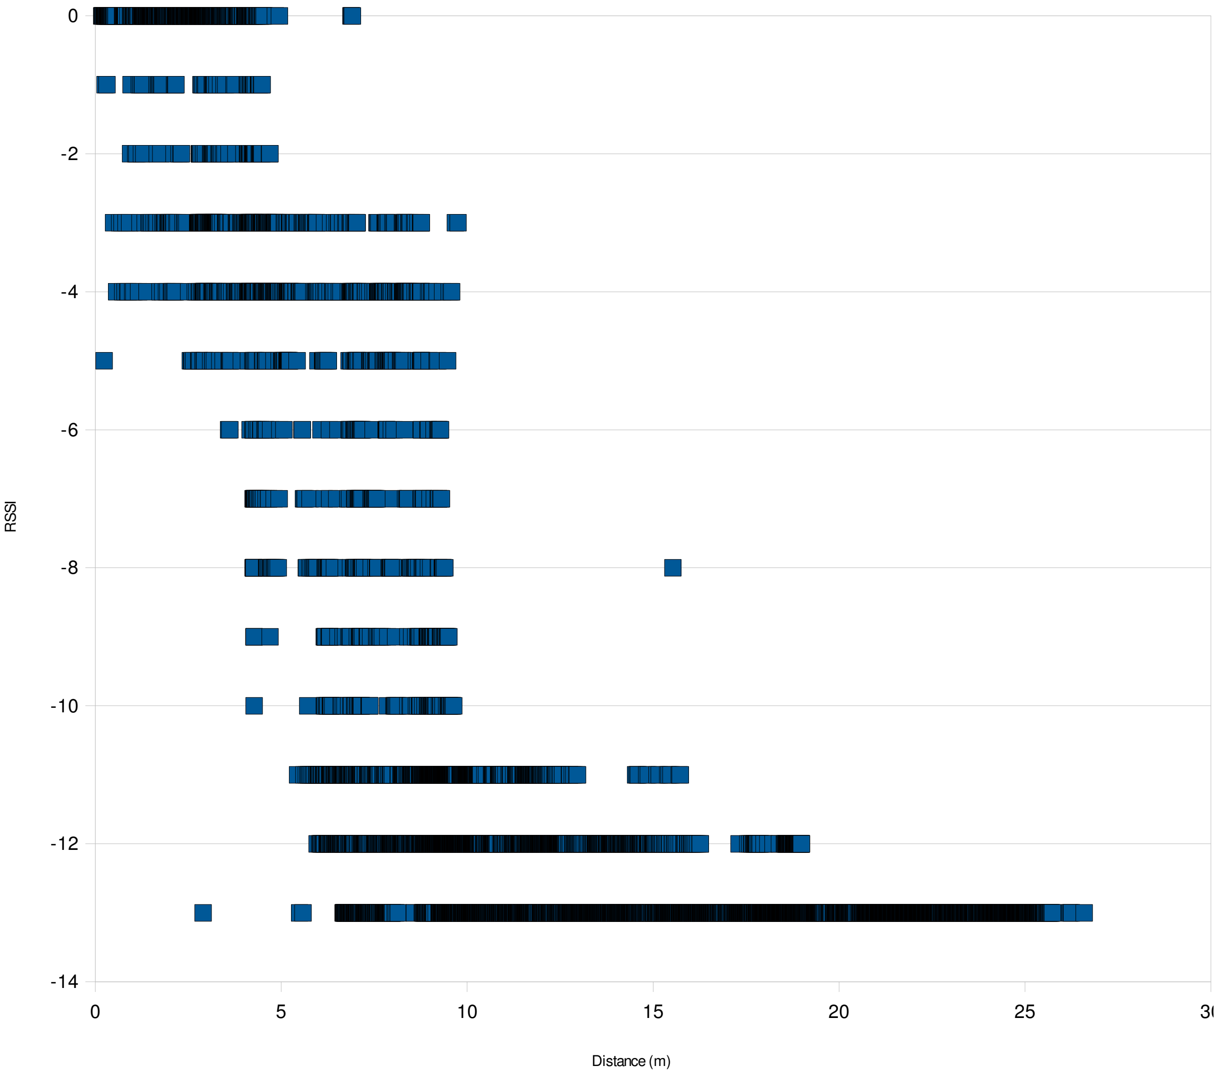

Figure 5.11: Measured RSSI against Euclidean distance from the host [87]

5.4.5.1 System architecture

The system architecture assumes that a large number of inexpensive Class 2 Bluetooth radios are distributed throughout a building. Each radio must be connected to a host machine of fixed, known location and the host machines must be networked together in some way. In most scenarios, this would entail at most simply adding a Bluetooth dongle to desktop PCs, making deployment both quick and easy; in fact, many new computers ship with Bluetooth adaptors built in. These fixed machines with Bluetooth are referred to simply as *hosts*. Their locations can be obtained directly from OpenRoomMap (Section 2.5.2), removing the need for any additional manual calibration or data entry to benefit from indoor tracking.

A central system that can connect to any host maintains a database of handset device IDs, their owners and any known properties of the handset (T_{pg_scan} etc). Each entry in the database is the result of a one-off registration of the handset with the system.

The central system then maintains a model of where people are and uses this to instruct hosts to monitor associated handsets at appropriate times. For example, as a user left their office the system may instruct hosts at either end of the corridor to monitor the handset in order to infer the direction in which he or she went.

5.4.5.2 Distributed clock offsets

This system involves continually connecting and disconnecting hosts and handsets and so it is extremely important to minimise the connection times. From the analysis of Section 5.4.2, ideally the host always incorporates the handset's current listening frequency within the paging train A. If only one host was involved, this could be achieved by accepting a lengthy first connection and thereafter caching the clock offset associated with that handset.

However, the connection/disconnection here involves many hosts, typically in sequence. Once a host discovers the offset between its local Bluetooth clock and that of a handset, it should distribute the offset to other hosts. For this to work, each host must be able to estimate the offset between its local Bluetooth clock and that of the handset. The estimate must be sufficiently accurate that the host predicts a listening frequency for the handset that is within 8 frequency hops of the true value. This ensures that the true frequency will be in train A.

The handset advances its listening frequency every 1.28 s. Thus, hosts other than the original may incorporate an error of up to $8 \times 1.28 = 10.24$ s in their clock offset and still connect using train A. Therefore distributing clock offsets can be trivially achieved by using NTP to synchronise the system clocks of all hosts, and having them report two offsets: the clock offset between their system clock and their Bluetooth clock, and the offset between their Bluetooth clock and that of the handset.

To verify this approach in principle, two desktop machines were time-synchronised using NTP. Both machines in turn were connected to a particular handset and the offsets between the local Bluetooth clock and both the handset clock and the local system clock were measured. It was possible to predict the offsets of one machine given only the data from the other to an accuracy far greater than was required to connect in train A. In practice, however, the BlueZ code that permits specifying the clock offset to use for a particular connection is not yet functional. It is therefore justifiable to assume that all connections take place within T_{pg_scan} seconds, which gives an expected connection time of 0.64 seconds in the default case.

5.4.5.3 Search

At its simplest level, this setup can be used to emulate an scan-based system for a single user. A subset of hosts that represents the minimum set needed to cover the entire building would simultaneously be instructed to connect to the associated handset. Any successful hosts then report back, allowing the central system to localise the user. Those devices that are out of range and still attempting to connect after localisation has completed can simply cancel their paging, allowing the user to be localised quickly and the next update to begin.

Whilst this algorithm will localise the user significantly faster than an equivalent scan-based system, it does not scale well with the number of users. An scan-based system may only provide position updates every 10.24 s, but it handles multiple handsets simultaneously. A connection-based system used in the manner described should be able to localise around ten handsets in the time that an scan-based handset would localise almost all handsets.

This shortcoming can be addressed in two ways. The first is to use the connection monitoring techniques of Section 5.4.4 to identify those users who are moving away from their nearest host (the current 'home' host) and who therefore need to be tracked. From previous analysis of people's working habits it is known that the vast majority of workers will be sedentary at any given time and therefore the number of mobile users should be small enough to localise all users within reasonable time-frames [89].

The second is to limit the extent of search where possible. If, for example, a user's current home reported a very weak RSSI (indicating lost proximity) t seconds ago, the only hosts that can reasonably observe him now are those within approximately vt of the reporting host, where v is the speed of the user. This type of information permits simultaneous searching for different handsets known to be in different building areas. This, combined with the fact there may be disjoint subsets of hosts that completely cover any given region, should permit tracking of many users simultaneously with update rates similar to that of scan-based tracking. However, the system will struggle to cope with many users moving at once or collected within a small area with few covering hosts. In such a scenario, the update rate for tracking will degrade, but the system will self-heal when the crowd disperses.

5.4.5.4 Bootstrapping

The discussion of TcB systems so far has ignored how a registered handset returning to the building will be discovered before being tracked. There are currently two options:

Constant round-robin polling. The obvious solution is to cycle continually through all the registered handsets that are not being tracked and poll them from a variety of hosts. Because a connection can take up to 5.12 s to connect (worst case), cycling through 100 registered handsets can take over eight minutes to complete. Additionally, it places an extra load on hosts that may be in use for monitoring tracked handsets, reducing update rate.

Out-of-band events. The preferred solution is to use out-of-band events to signal that a particular handset is likely to be in the building. These events can be easily generated from computer system events (e.g. the associated user logs in) and from building security (e.g. the user gains access to the building by swiping their ID card). It may also be possible to infer presence based on cell ID from the mobile telephony network.

Additionally, many people follow a daily routine which there is potential to learn autonomously. Doing so could allow prediction of when a user is likely to arrive, permitting for targeted polling in future of the subset of expected handsets rather than all handsets.

5.4.6 A target-connects-base tracking system

The predictive approaches to connection-based tracking described so far work well for small systems, but the need to bootstrap ultimately limits scalability. In a BcT system without the benefit of out-of-band events to bootstrap, each base must continuously scan for all possible targets in case a new target appears; in a TcB system the need to initially search for the full set of bases slows the technique.

One remedy is to make use of *spoofed* Bluetooth addresses, as mentioned in Section 5.4.1.1. Some Bluetooth platforms allow for a Bluetooth device to temporarily and arbitrarily alter its BDADDR.[10] A TcB system can exploit this by using a network of bases each with the same (spoofed) address. Targets then continually attempt connections to that address, exchanging data with any successfully paged device in order to identify it. For this technique to work reliably the paging ranges for each device should not overlap. If they do, there is a chance for reply collision, with multiple bases responding to the same page. Such collisions greatly confuse some Bluetooth stacks, and it is best to place the spoofed bases such that they do not have overlapping coverage.

It is possible to invert this system to form a spoofed BcT system where the targets all adopt the same address, which the bases duly page. When multiple targets are within range of a base, however, their page replies may collide. Therefore this system is only appropriate when the targets are expected to be spatially disparate at all times. This is an unlikely scenario and therefore not considered further here.

5.5 Security and privacy

5.5.1 Location privacy

Many people are sensitive about divulging their whereabouts automatically, with good reason. There are three separate privacy concerns: privacy from observers external to the system (*attacker privacy*), privacy from components of the location system itself (*system privacy*) and access control to dynamically opt in or out of tracking altogether.

5.5.1.1 Attacker privacy

Ideally a system will have high attacker privacy, meaning devices that are not part of the tracking infrastructure cannot track the target. In practice this is difficult to achieve. This discussion assumes that the attacker has no access to the location system as a whole, but can plant Bluetooth devices of their own at will.

A TsB system offers good attacker privacy since the attacker cannot identify the origin of an inquiry. Conversely, the BsT system requires that the target be discoverable, meaning it is discoverable to the attacker too.

The BcT technique has the significant advantage that target can remain undiscoverable. There is, however, an important caveat: in the paging process the pager must identify both the device it is paging and itself. This must be done openly since otherwise the listening device would not know which decryption key to use. In principle, a Bluetooth sniffer could detect any paging messages exchanged and hence identify the target. However, such sniffing equipment is expensive, bulky and highly specialist and so this risk is small.

[10] See the `bdaddr` tool in the BlueZ source code.

If the TcB approaches make use of address spoofing, an attacker's device could adopt the same address and act as a base. The target would then connect to the fake base, revealing its presence. Whilst this cannot be prevented, it can be detected if the target periodically compares connection logs with the system.

5.5.1.2 System privacy

The Bluetooth specification has the property that any device initiating an inquiry *cannot* be identified by those that receive the request. In a TsB system, the bases know they are being scanned, but not the device that is scanning. The target can therefore decide whether or not to disseminate its location information and TsB has good system privacy. In BsT, the bases must be able to identify the target and hence the location information is gathered by the system and not the target, resulting in poor system privacy.

The two connection-based techniques also have weakened system privacy, since each end-point of a connection must know the address of the other. Address spoofing may allow a target to adopt a pseudonym address that it changes regularly.

5.5.1.3 Access control

It is also important that there is some degree of access control for the target device—a user should be able to opt out of tracking at any time. By default most Bluetooth devices enable the page scan mode, by which they periodically look for pages as described previously. When in this state, a BcT approach will be able to identify the presence of the target. However, the attacker would need to know the native address of the target for this to work (assuming the target has no reason to spoof its address in general), and tracking multiple devices is difficult.

Hence the best technique to limit tracking of a device is to turn off the transceiver altogether. Although this makes for a clear tracking indicator to the user, it does limit normal Bluetooth functionality.

5.5.2 Security

Since the mobile targets for personal energy metering are likely to be personal devices, any vulnerabilities introduced by the Bluetooth tracking modes must be assessed for risk. The most obvious security risk is for the target to advertise its presence through discoverability, as per BsT. With inquiry scanning permanently enabled attackers can quickly find the device and exploit any vulnerabilities in the Bluetooth stack in use to gain unauthorised access to the device. Even without such vulnerabilities an attacker can reduce the battery life of the target by continuously performing inquiries on it. Worse, because inquiry responses do not raise any HCI events in the inquired device and the inquirer does not identify itself, it is generally not possible to observe such an attack. In the course of gathering the power measurements of Section 5.7, a rogue device performing a continuous scan was inferred from observed regular power spikes and it was necessary to search for and turn off all devices before the rogue device was identified.

For the connection-based techniques there is only the potential for an attacker to snoop at a low level using specialist equipment. This would allow them to identify a nearby device similar to it being discoverable. However, a standard Bluetooth device could not be used for this snooping without extensive modification.

5.6 Tracking evaluation

Working with Bluetooth can be a frustrating experience since different manufacturers may not have a fully consistent interpretation of the specification, different stack implementations support different subsets of the specification and stack stability is still improving. Consequently Bluetooth systems should never be analysed purely theoretically. This section demonstrates a number of tracking systems formed using the techniques described above. Although quantitative results are very dependent on a number of factors including the geometry of the base deployment, the layout and construction of the tracking space and the path and time taken during the trials, they nevertheless reveal a lot about the performance and utility of the systems.

Although each of the systems has use cases where it will be most appropriate and the personal energy meter will rely on input from heterogeneous systems deployed in different buildings, in general TcB and TsB are most likely to scale well to large areas with many users where no additional bootstrapping information is available. These two systems are therefore evaluated further in this section.

5.6.1 TsB tracking evaluation

To test a TsB system seven bases were positioned throughout an office corridor as shown in Figure 5.12(a). As before, the situation was as representative as possible: Bluetooth had to co-exist with WiFi, walls attenuated signals and those not participating continued to work as usual. Each base was set to discoverable with the following parameters (quoted hereafter in 625 μs slots):

$T_{w\_inq\_scan}$	4,000
T_{inq_scan}	4,096

In effect, this meant that each base was almost continuously listening for inquiries. This prevents them from having meaningful Bluetooth connections themselves, but means response rates are fast and that the target therefore saves energy by not having to transmit inquiry packets for as long. To have a chance of discovering all the bases, the target needs to inquire until at least the first train change occurs (2.56 s for the specification-required $N_{page} = 256$). In the test each scan was run for 5.12 s.

A walking test was performed with a laptop set to continuously scan and to collect all HCI inquiry responses. The ground truth location of the device was determined using a Bat attached to the laptop user.

Figure 5.12 illustrates the results for each of the seven bases. The thin purple line shows the route taken by the user, whilst the thicker blue lines indicate parts of the walk where inquiry events for the relevant base were being received. In practice, however, a target cannot know that it will not receive any further inquiry events from that base until a complete scan has occurred without an event. The thicker green lines show the parts of the walk that a real target would have to consider itself in range of the base.

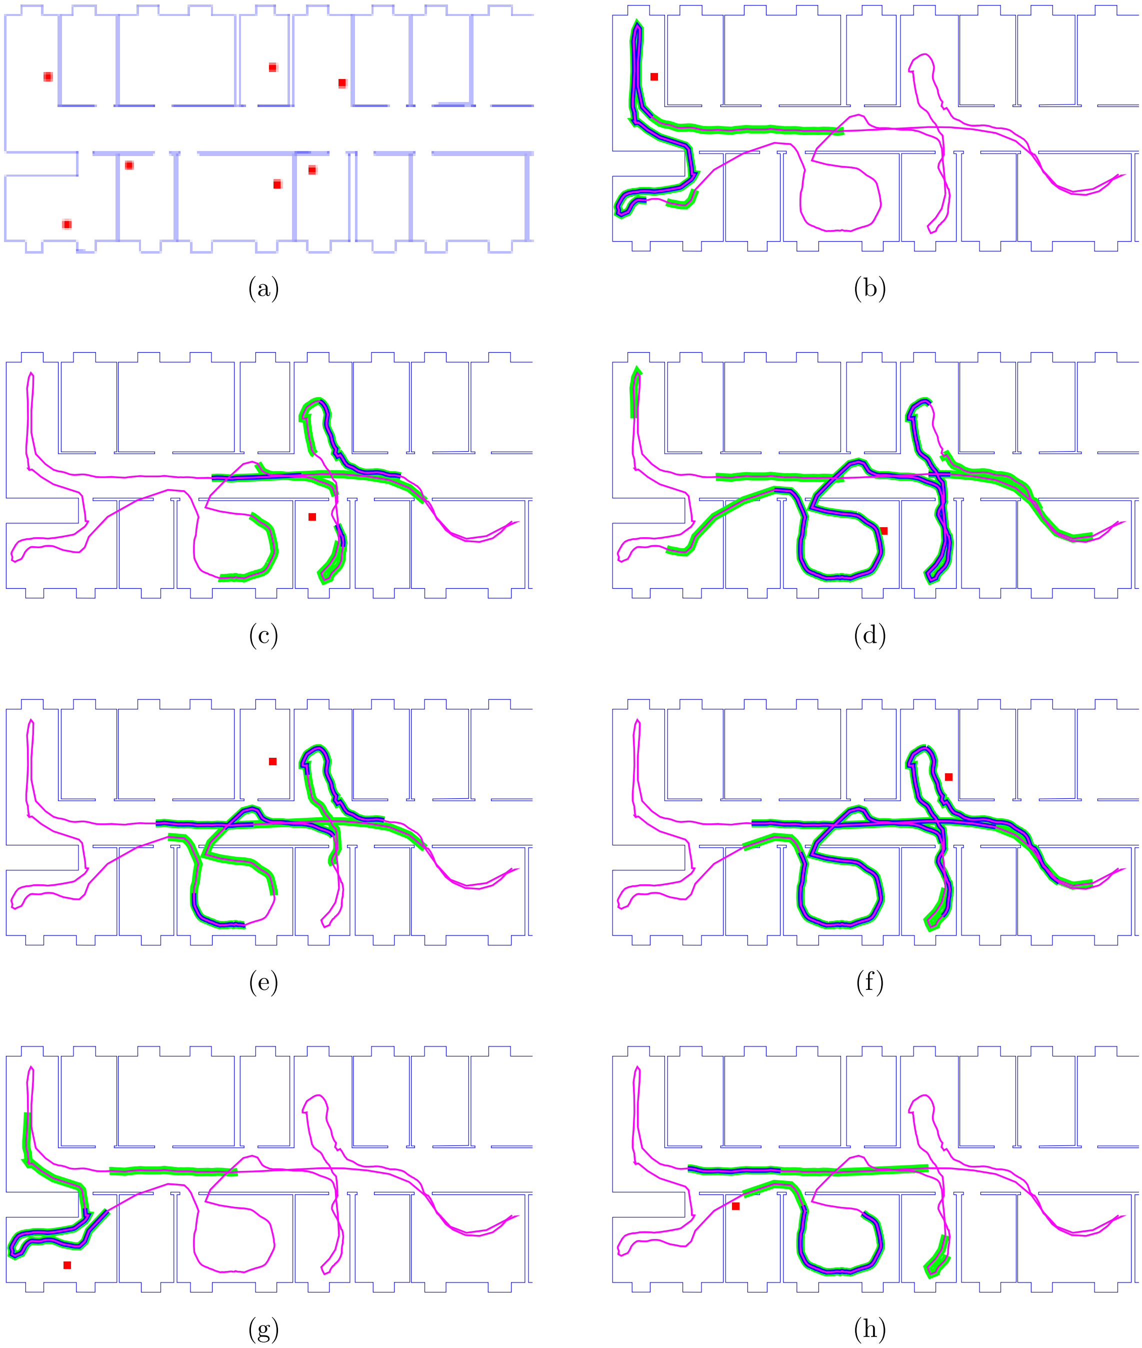

Figure 5.12: TsB system results. (a) Positions of the seven bases. (b–h) Tracking results for each base.

5.6.2 TcB tracking evaluation

The most scalable TcB mode involves the use of bases with a common spoofed address. This technique was applied to the arrangement of bases shown in Figure 5.12(a), all of which were set to have the same arbitrary address. As is clear from the TsB results, the bases had overlapping coverage, which is problematic when multiple devices assume the same address. To limit the paging area for each base EMI shielding tape was applied to the Bluetooth adaptor in the laptop. Two layers of 3M 1345 EMI shielding tape were sufficient to reduce the range such that overlap was minimal. This tape could instead be applied to the base adaptors if the target does not have an exposed Bluetooth antenna (e.g. on a mobile phone). The shielding increases granularity at the cost of requiring an increased number of base stations to ensure coverage; the appropriate tradeoff will vary depending on the level of energy apportionment required. If building-level apportionment is sufficient (perhaps because no more fine-grained consumption data is available), Class 1 Bluetooth devices with a nominal range of 100 m could be used instead, reducing further the deployment complexity.

A similar test to that for TsB was then performed, setting the base paging parameters to be:

$T_{w\_pg\_scan}$	18
T_{pg_scan}	2,048

These parameters were chosen because higher duty cycles have two disadvantages in a TcB system using spoofing. Firstly, where coverage areas do overlap, it greatly increases the chance of collisions when multiple bases respond to the same page. To see this the target was placed near to one base, but potentially in range of three others with the same spoofed address. It was then left (without EMI shielding) performing remote name requests on the spoofed address, varying the duty cycle. A failure rate of 0.38% was observed with a $T_{w\_pg\_scan}$:T_{pg_scan} of 18:2,048, which rose to 4.01% using 18:512. With the shielding tape applied, only the adjacent base was sighted, and the failure rate was only 0.003% using 18:2,048.

The second disadvantage of a high duty cycle for connection-based tracking is that it requires the base radios to be almost constantly listening and therefore unable to maintain even the remote name request connections long enough for reliable communication. Neither disadvantage is present for the TsB system because it does not use spoofing nor does it need to maintain connections.

The particular settings used were trialled because they are the default BlueZ settings. The results were sufficiently good that there was no need to alter the duty cycle further, although other system implementers could optimise the parameters for their particular usage scenario.

Figure 5.13 illustrates the test results, which were collected using remote name requests from the laptop to the special address. A coloured line is drawn from each base to the position of the user at the moment the HCI layer reported the remote name. Black crosses are used to indicate the report of a name request failure. There is a strong spatial locality for each of the bases, and a good update rate when in range of the bases. The connection failures can be attributed to both unintentional areas of overlapping coverage and coverage holes from an unoptimised distribution of bases. Even without these optimisations, as would probably be the case in the majority of deployments intended solely for energy metering, the connection-based approach provided results comparable to the scan-based equivalent.

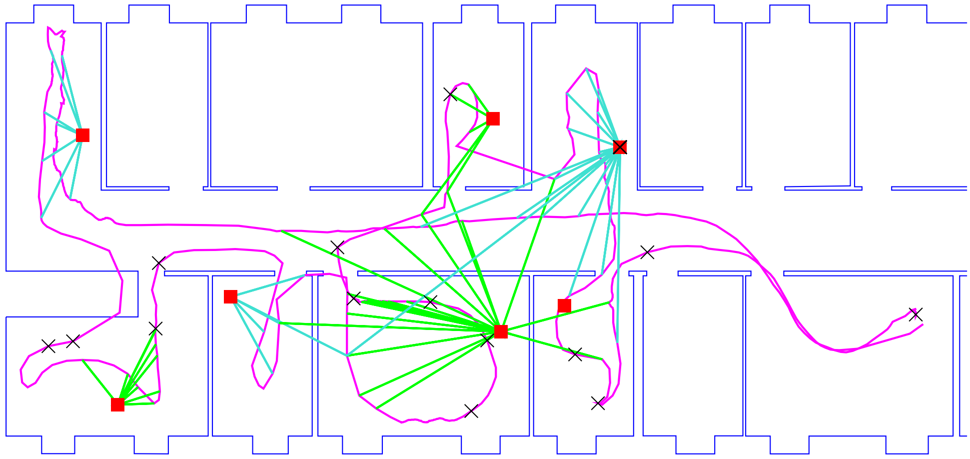

Figure 5.13: TcB system results.

5.7 Battery costs

5.7.1 Energy proportional location systems

There is little point building sensor systems with a view to reducing global energy consumption if the energy used by those systems outweighs the possible savings. Instead, we must strive for an *optimal digital infrastructure* which is implemented in energy-efficient ways and is operational only when delivering a service for some real end-use [97].

To date, very little consideration has been given to the energy consumption of location systems. For example, the Bat system relies on a network of ultrasound receivers installed in the ceiling [3]. These receivers are permanently on, drawing approximately 25 W per room; assuming a personal load of 150 W (Section 3.5.5) and an 8 hour working day, the Bat receivers in a single office account for half as much energy as the rest of the devices.

Some attempts have been made to reduce the infrastructure required by a location system—in particular, Jevring et al. demonstrated dynamic optimisation of their Bluetooth localisation network [103] and Nishida et al. investigated the number of receivers required in ultrasound-based systems [150]. These are generally aimed at simplifying administration rather than reducing energy consumption. In the Bluetooth-based systems proposed here no additional infrastructure is required, but battery power is always a key issue for mobile devices. Providing any new functionality such as tracking ability is useless if users disable it to save power. This section therefore quantifies the energy costs of using Bluetooth in this way and so demonstrates the capabilities of the profiling framework described in Section 4.5.

A key design decision for Bluetooth was that it must have low power consumption; this is also a reason for its choice to underpin location services. The actual power drawn will be dependent on many factors, including the Bluetooth chip in use, the antenna setup and the Bluetooth stack. However, the general trends are revealing and results are shown for several modern smart phones.

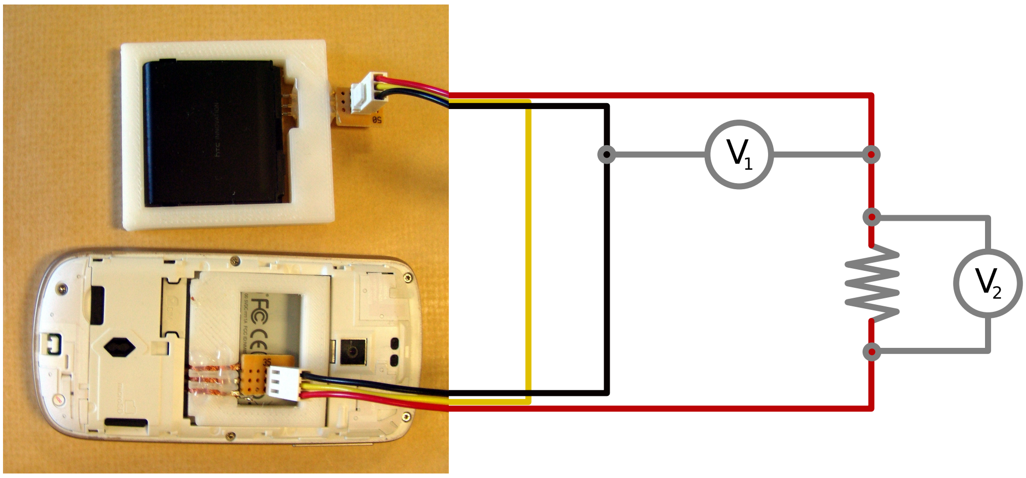

Figure 5.14: Replacement battery and battery holder for Magic handset.

5.7.2 Measuring mobile phone consumption

There has been little work on in-depth energy monitoring on mobile phones. Platform manufacturers have offered some tools and guidance, but users have traditionally only cared to know the total energy remaining. As a result, most systems, such as Nokia's Energy Profiler,[11] tend to have low resolution with slow response times.

It is very difficult to model accurately the energy cost of the components of a generalised handset. Instead, the technique described in Section 4.5 for decomposing power measurements allows accurate sampling of power draw specific.

The power consumption of mobile devices can be measured by replacing the battery with a printed plastic replacement[12] and inserting a high-precision 0.02 Ω measurement resistor in series between a battery terminal and its connector on the device (Figure 5.14). Replacement batteries and battery holders were produced to fit the G1,[13] Magic,[14] Hero[15] and Nexus One[16] handsets (Figure 5.14). In all of the tests that follow the G1 and Magic handsets (running Android version 1.1) produced indistinguishable results. The results for the Hero handset are for Android 1.5 and the Nexus are for Android 2.1.

A National Instruments PCI-MIO-16E-4 sampling board measured the voltage across the phone battery and also the voltage drop across the measurement resistor (which is first amplified with an instrumentation amplifier) at 250 kHz. Inserting the measurement resistor increases the circuit resistance, and therefore its power consumption. This is not a problem for these purposes as the increase is typically less than 1% of the total power. The measurement points are shown in Figure 5.14; simple rearrangement and application of Ohm's law yields $P \propto V_1.V_2$.

The presence of high-frequency components within electronics does not cause exceptionally rapid changes in the power consumption. This is most likely to be due to buffering within the device caused by capacitance in the circuit or voltage regulation. Inspection

[11] `http://www.forum.nokia.com/Library/Tools_and_downloads/Other/`
[12] The replacements were produced using a Reprap 3D printer (`http://reprap.org/`)
[13] `http://www.t-mobileg1.com/`
[14] `http://www.htc.com/www/product/magic/overview.html`
[15] `http://www.htc.com/www/product/hero/overview.html`
[16] `http://www.google.com/phone`

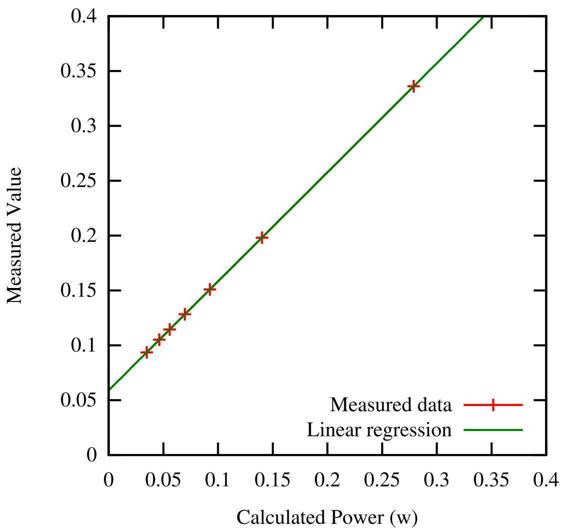

Figure 5.15: Calibration with known resistance

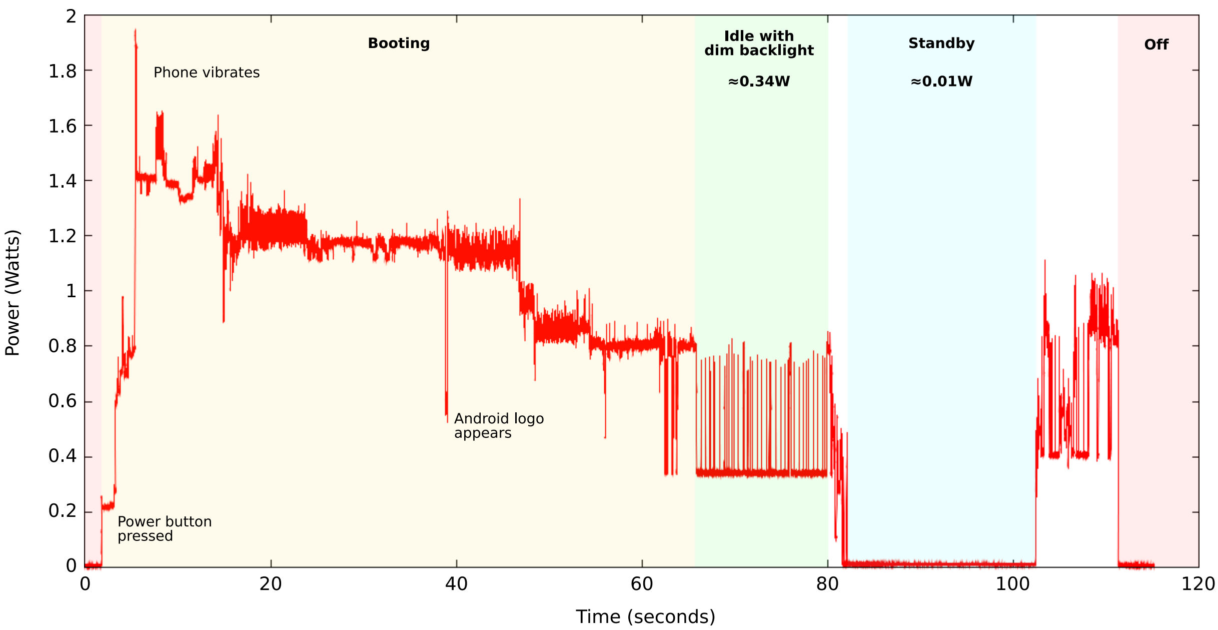

Figure 5.16: Instantaneous power consumption when switching on and off a G1 mobile phone [176]

of a number of device traces using a high-speed (1 GHz) storage oscilloscope confirmed that the sampling rate of 250 kHz was sufficient to capture all features of the trace.

The expected power consumption of a resistor is trivially calculated and so a selection of high-precision resistors can be used to calibrate the device. Figure 5.15 shows on the y axis the measured power draw and on the x axis the power calculated using Ohm's Law. The necessary scale factor was then computed using a linear regression through the resulting data points, with an RMS of residuals of 0.0001 (4dp).

Figure 5.16 shows an example trace when switching on a G1 mobile phone. The annotations of each boot phase and average power measurements were made manually.

Test ID	Handset State	G1 (mW)	N80 (mW)
1	Idle	12.81	19.44
2	WiFi connected but idle	170.66	438.08
3	Bluetooth on, discoverable, unconnected	15.45	22.17
4	Bluetooth on, continually being scanned by host	16.07	31.80
5	Continuous Bluetooth echo at maximum rate	321.07	234.43
6	Continuous Bluetooth RSSI measurement	74.97	89.20
7	Bluetooth echo every 30s (with reconnect)	23.76	26.17
8	Bluetooth echo every 20s (with reconnect)	25.19	28.17
9	Bluetooth echo every 15s (with reconnect)	28.02	–
10	Bluetooth echo every 10s (with reconnect)	30.53	42.27
11	Bluetooth echo every 5s (with reconnect)	40.13	50.61
12	Bluetooth RSSI every 30s (with reconnect)	29.93	28.05
13	Bluetooth RSSI every every 20s (with reconnect)	35.86	29.93
14	Bluetooth RSSI every every 10s (with reconnect)	47.59	36.04
15	Bluetooth RSSI every every 5s (with reconnect)	75.72	51.88

Table 5.4: Experimentally measured power draws

5.7.3 Connection monitoring costs

The hardware described above was used for a preliminary investigation into the power draw associated with monitoring a connection using either RSSI or echo response times on the G1 and Nokia N80 handsets. These would be the costs associated with a BcT tracking system.

The results are reported in Table 5.4, with each test given a unique ID for reference here. Tests 1–3 provide baseline power draws; tests 4–6 provide power draws for continuous connection states; tests 7–15 provide power draws for discontinuous state (e.g. the connection was created, an RSSI taken and the connection shut down until the next update time). For all of these tests, the handset screen was powered off, no SIM card was in place, no other connections were active and no applications (apart from any operating system requirements) were running on the devices. Bluetooth parameters were left at their defaults. The power draws were computed by averaging over appropriate time periods (mostly 10 minutes). Note the very high cost of WiFi compared to Bluetooth, highlighting one reason it is less suitable for the sort of continuous tracking required for energy metering.

The results are broadly as expected. A handset that is page scanning need only take action every T_{pg_scan} seconds, and then only for 11.25 ms, so only a small associated cost is observed. Saturating a permanent connection with echo requests (test 5) is very energy-intensive, whilst continually measuring its RSSI lies between the two extremes (and equates to the energy cost of maintaining an L2CAP connection).

When considering less frequent monitoring, instead of holding a connection open it may be more energy efficient to connect, take the reading, and then disconnect. Tests 7–15 concerned connection monitoring with such a model and, within error bounds, there is

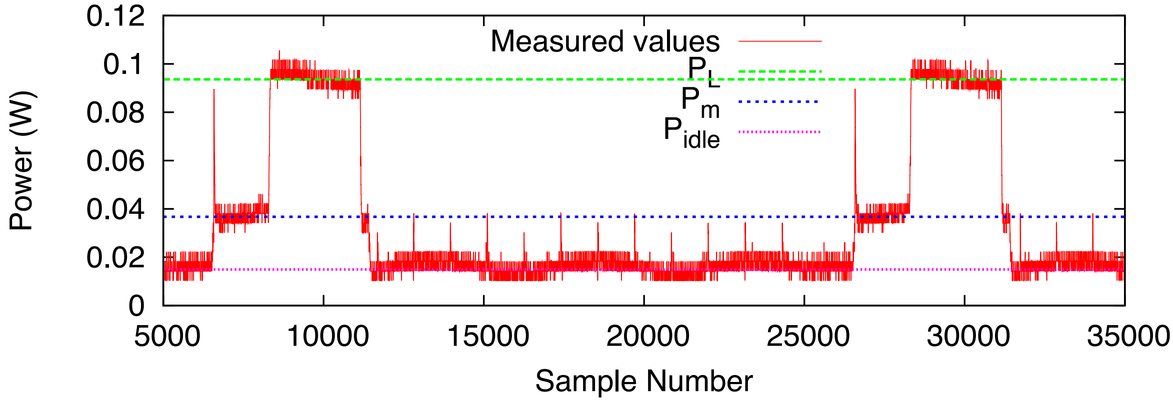

(a)

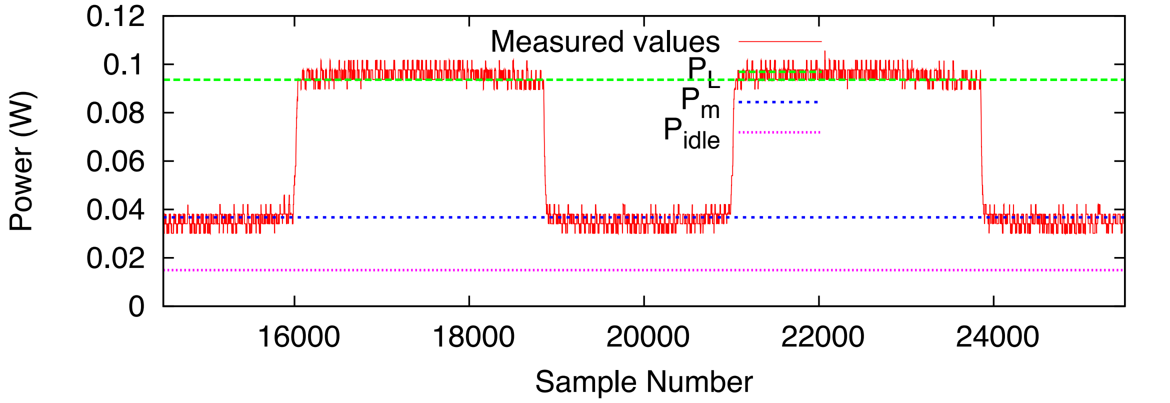

(b)

Figure 5.17: Typical power traces for different Bluetooth settings, sampled at 250 kHz. (a) $T_{w\_pg\_scan}$ =18, T_{pg_scan} =128. (b) $T_{w\_pg\_scan}$ =18, T_{pg_scan} =32. Note the absence of the idle power state in (b).

little to choose between using an echo packet or reading an RSSI. The main cost lies in creating and destroying a connection, which is relevant to both metrics. These data show that if the update period exceeds 5 s, it is more efficient to break the connection and reform it when it is next needed than maintain the connection. For period less than 5 s it is unlikely to be possible to disconnect since the underlying ACL does not close with the L2CAP connection.

To understand better the costs of listening states the energy profile of the device should be related to the listening parameters varied in Section 5.3.1. A set of experiments was therefore carried out using a Nexus One running Android 2.2 and rooted to provide the ability to change BlueZ parameters. The data collection was performed with the phone inside a shielded box to prevent external radio influences and where possible the phone was idle (i.e. in standby) since it is in this mode that phones are expected to spend the majority of their time, and there are then fewer events in the power trace for which to account.

Figure 5.18 shows the variation in power consumed with the $T_{w\_pg\_scan}$ parameter for the Nexus One whilst T_{pg_scan} =18. Closer inspection revealed that the smartphone exhibited at least three relevant power states as illustrated in Figure 5.17: a power P_L associated

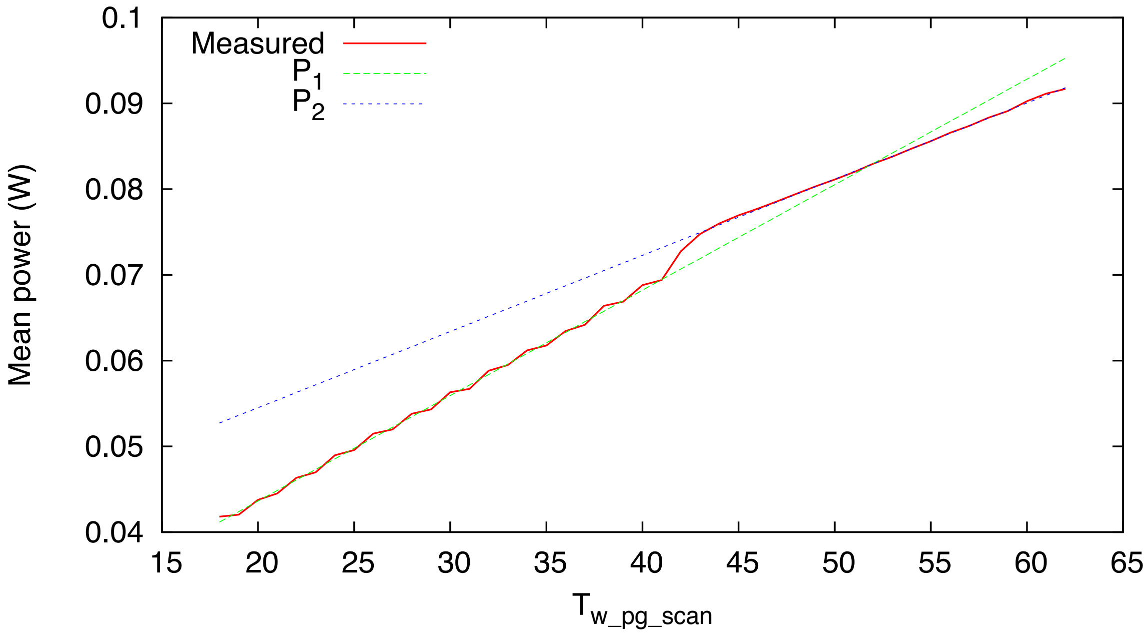

Figure 5.18: Variation in Bluetooth power consumed during inquiry listening

with the listening state; a power P_m drawn for a constant period of time (T_m slots) prior to starting the listening window; and a power P_{idle} associated with the standby power of the smartphone. This was modelled as:

$$P_1 = \frac{T_{w\_pg\_scan} P_L + T_m P_m + (T_{pg\_scan} - T_{w\_pg\_scan} - T_m) P_{idle}}{T_{pg\_scan}} \tag{5.3}$$

where P_L =0.0936 W; P_m =0.03674 W; P_{idle} =0.0149 W; T_m =12.112. However, for high duty cycles, the phone never returned to the idle state and the output is better modelled by:

$$P_2 = \frac{T_{w\_pg\_scan} P_L + (T_{pg\_scan} - T_{w\_pg\_scan}) P_m}{T_{pg\_scan}} \tag{5.4}$$

Figure 5.18 overlays these models on the measured data. The equivalent experiments were repeated for inquiry scanning (whilst page scanning was turned off) and the analogous models were found to apply with the same values for the constants. This was expected since listening for a page packet is no different from listening for an inquiry packet. Note, however, that the phone appeared to treat inquiry scanning and page scanning as completely separate processes, causing a mix of the models when both were enabled.

In general, a duty cycle greater than approximately 65% precipitated a switch from P_1 to P_2. In the context of tracking, these data show that a linear relationship between the duty cycle and the power consumption is a reasonable assumption. As a final note, the power consumed whilst in standby and scanning for inquiries with $T_{w\_pg\_scan}$ =18 and T_{pg_scan} =32 was measured at 0.0695 W—three times greater than the normal standby consumption of the Nexus One.

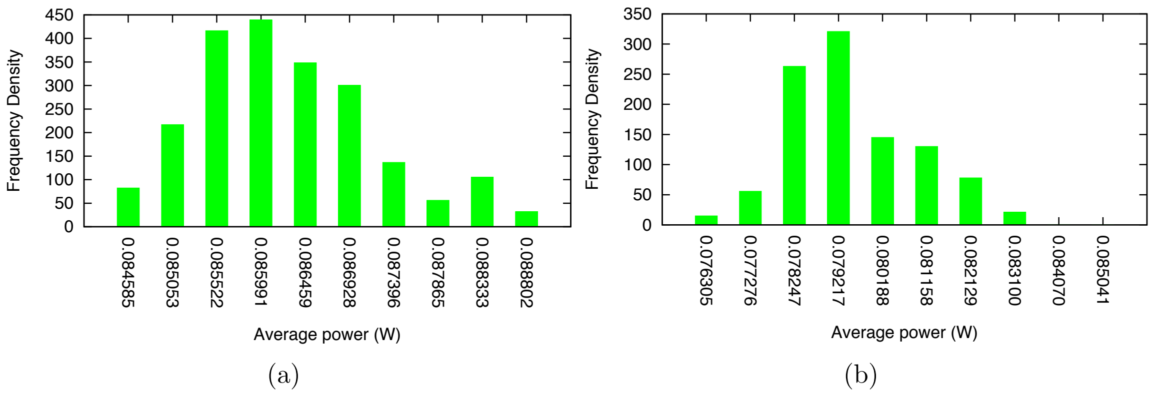

Figure 5.19: Distribution of average power drawn over 5.12 s for (a) inquiring (b) paging.

5.7.4 Cost of paging or inquiring

For a device in the paging or inquiry substate the Bluetooth radio will be constantly sending data and so one would expect a roughly constant power draw. Figure 5.19 shows the distribution of average power drawn whilst paging a device that was not present. Each power value was an average of the measured powers over 8,192 slot (5.12 s) timeouts. Power draws of around 0.08 W were observed—to put this in context, the idle power consumption of the phone was approximately 0.02 W.

5.8 Discussion

Connection-based tracking is a viable alternative to scan-based tracking, although neither is ideal. The advantages of connection-based tracking as:

Faster update rate. Connections can generally be *established* faster than a scan can be completed. Connections can be *monitored* at a very fast rate.

Variable update rate. The rate can be adapted to suit the context of the mobile device and preserve battery life. If, for example, the user is believed to be sedentary, the query rate can be reduced.

No scanning. The lack of need for constant inquiries is a significant enhancement to Bluetooth usage, security and privacy.

Privacy. Users retain their right to opt out of being tracked by switching off Bluetooth on their handset.

There are, however, a number of drawbacks:

Devices must be known. Connection establishment requires that at least one party knows of the existence of the other in order that it can issue a connection request; however, this convenience drawback is also a privacy advantage.

Security policies are not universal. Not all Bluetooth security policies permit L2CAP connections without authorisation. Similarly, not all Bluetooth chips permit querying the connection metrics such as RSSI. However, such queries are performed on the beacon's Bluetooth chip, which can be selected accordingly.

Connection limits. Care must be taken to ensure the Bluetooth limit on the number of concurrent connections (usually 7) is not reached; this limit can reduce the update rate in a BcT system with many nearby targets since, once it is reached, a base cannot attempt to connect to one target until it disconnects from another.

Figure 5.5 compares connection and inquiry based tracking side by side. Ultimately, Bluetooth was designed as a short-range wireless communications protocol and as such it is unreasonable to expect it to provide ideal location tracking. Nonetheless, both scan-based and connection-based tracking are at least feasible, and both can be implemented without modifying handset software or hardware, which is an important consideration.

Note that:

- Scan-based tracking can provide higher update rates than usually quoted in the literature. The use of HCI events and manipulation of the inquiry parameters can provide many RSSI measurements per inquiry, and each inquiry need only last a few seconds rather than the full 10.24 s.

- Bluetooth address spoofing is a viable option for connection-based tracking, and significantly enhances scalability. Paging parameters can be optimised to speed up paging times. The remote name request is a more robust technique for establishing an short-lifetime ACL.

- Although Bluetooth was designed for low power consumption, the choice of page and inquiry parameters is important; while sensible settings would allow continuous tracking throughout the delay with minimal effect on battery life, in the worst case, standby power consumption can be three times greater than normal.

Looking forward, connection-based tracking will improve in the short-term as handsets advance. Already handsets incorporate accelerometers (which could be used to infer movement or the lack of it, or to dynamically update the monitoring rate). Similarly, advanced development platforms for mobile devices are emerging and these will hopefully provide low-level access to handset subsystems such as Bluetooth. This in turn will permit handsets to be more active in the location process since APIs and behaviours will standardise and applications will be easier to deploy.

5.9 Summary

Although a personal energy meter will rely on a heterogeneous network of location systems, few existing systems are suitable for widespread use due to the extensive bespoke infrastructure that must be installed and surveyed. This chapter has identified and evaluated an alternative in the form of Bluetooth tracking and presented a series of novel

Table 5.5: Comparison of scan-based and connection-based tracking

	TsB	BsT	TcB	BcT
Works on unmodified handsets	Most[a]	Some	Most	Almost all[b]
Requires custom handset software	Yes	Often[c]	Yes	No
Requires discoverable handset	No	Yes	No	No
Requires separate data channel	Yes	No	No	No
Location update frequency	Low–Med	Low–Med	Low–High	Low–High
Scalability	Simple	Simple	Med[d]	Complex[e]
Deployability	Med	High	High	High
User privacy	High	Low	Med	Med–High
Handset power drain	High	Low	High	Low

[a] High update rates only achievable if target API exposes HCI events reporting.
[b] Multiple connection types well supported on mobile devices.
[c] Many devices do not allow discoverability to remain permanently on. Higher update rates depend on ability to set target inquiry parameters—rare in embedded devices.
[d] May struggle with lots of targets in a small area. Address spoofing increases scalability. Multiple connection types so well supported.
[e] Scales poorly with no of targets

techniques, provided a detailed analysis of their strengths and weaknesses and demonstrated working tracking systems. The properties of Bluetooth that enable tracking have been studied both theoretically and experimentally, showing that there is potential for the approach. Although it is not the ideal tracking system, its use of existing infrastructure means it has very little in the way of deployment costs, making it a viable option for providing the context needed for a personal energy meter.

Chapter 6

Federation and scaling

Contents

Overview

This chapter examines a model personal energy meter to allow the techniques proposed in this dissertation to be scaled up to planetary proportions. It introduces a federalised, syndicated architecture, allowing allows independent systems to contribute energy 'feeds' to an aggregator that constructs a picture of an individual's energy use. It also presents a prototype implementation as a mobile phone application that visualises the data from sensor systems described in previous chapters.

6.1 Architecture

To piece together the complete picture of each person's energy usage requires information from a range of separate sensor systems which meter individual parts of his overall consumption. These systems may be widely distributed both geographically and in terms of their technical architecture; a mechanism is required to combine their outputs without the need for a centralised middleware layer and the ensuing questions about who should bear the responsibility and costs for its maintenance.

Inspiration can be taken from models that have proved successful in other fields. In particular, RSS and other syndication formats have proved popular on the web for publishing frequently updated works such as blog entries or news headlines in a standardised format.

David Piggott implemented the Android application described in Section 6.2.

A similar concept can be used to publish and subscribe to energy information; the easier it is to participate in the system the more likely it is to be widely adopted.

'Energy feeds' are analogous to the news feeds now pervasive on the web used to federate stories between sites. Instead of news, each feed contains periodic updates on the total energy used by a particular facility; users can subscribe to the feeds for facilities that are relevant to their lives. This is possible with only minor extensions to the Atom Syndication Format [151], meaning existing libraries and tools can be used and so easing adoption and deployment. Feeds contain details of the facilities they represent, including categories to aid aggregation. Each update includes a cumulative total energy consumption, meaning aggregators can check them as frequently or infrequently as is appropriate without fear of missing essential readings.

This decentralised approach presents a minimal barrier to entry: all that is required to participate in the personal energy meter network is to create a straightforward, text-based feed hosted somewhere on the Internet. Unlike web-services-based systems, this could be as simple as a static file; it could equally be the interface to a much more complex dynamic system, but in either case no central registration or approval is necessary. This simplicity compared to RESTful services [52] means feeds are easier to implement on resource-constrained embedded systems and can be cached elsewhere to combat problems with bandwidth or processing power. Furthermore, they allow historical data to be added when it is obtained; this is necessary for many of the systems described in this dissertation which rely on data not available in real time.

It also confers privacy advantages over centralised systems like Google PowerMeter or Microsoft Hohm discussed in Section 2.3.1.1: no significant infrastructure is required and no central party holds all the information about an individual. A user's complete energy profile could reveal a lot of information about him: where he went, what he did, who he associates with and what his habits are [82]. Instead, with a feed-based architecture, each sensor system knows only a small part of the larger picture; each refers to a person by a separate pseudonym.

6.2 Aggregation

Just as with RSS and Atom feeds of news on the web, there are several possibilities for the aggregation of energy feeds, including dedicated websites, 'widgets' embedded in other sites and client-side applications. There is also a 'hybrid' option such as that adopted by Google Reader[1], which provides a web-based interface to browse feeds but also offers an API to support a range of third-party desktop and offline clients. Different interfaces will suit different people, and a feed-based architecture allows each individual to select the aggregator that he prefers. Existing readers will be compatible, but it has been shown in numerous studies that the representation of feedback has a significant impact on its efficacy (Section 2.1). Each aggregator may therefore choose a different visualisation or feedback strategy.

To demonstrate the concept, an example energy aggregator was implemented as an application for Android-based mobile phones which presents both an energy stack and graphs

[1] `http://google.com/reader`

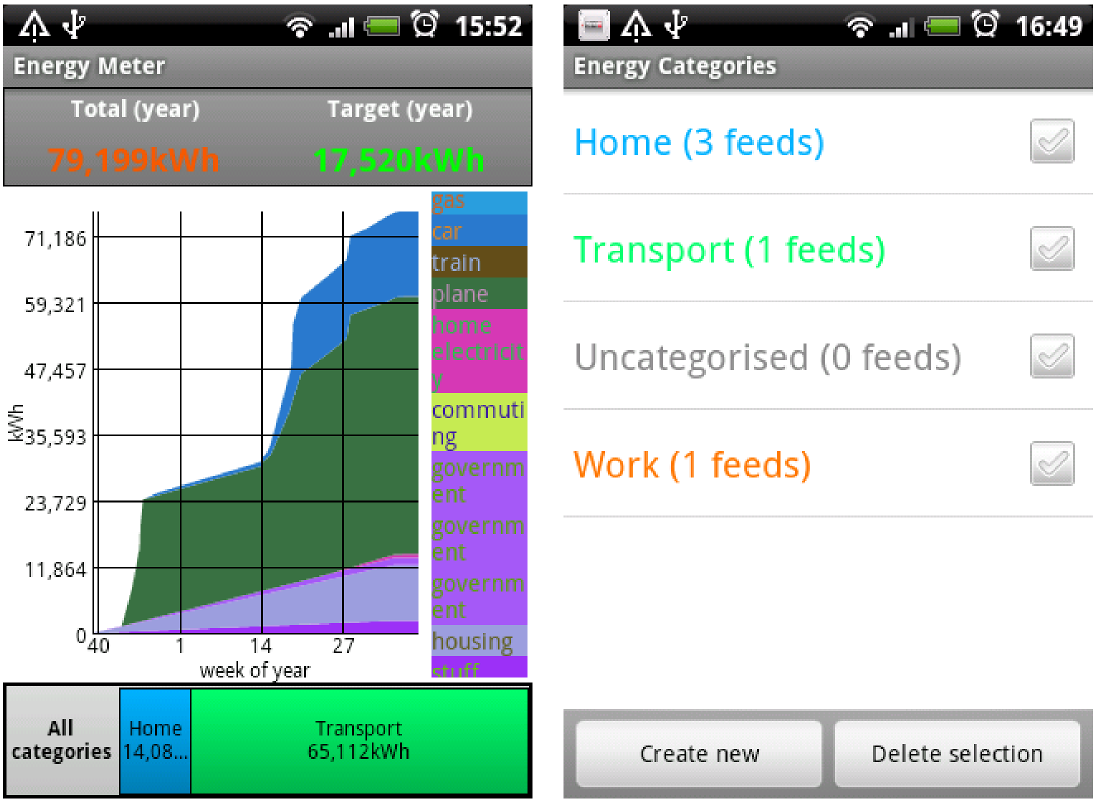

(a) cumulative totals and breakdown

(b) facilities are categorised

Figure 6.1: The Energy Meter application for Android-based mobile phones

showing cumulative consumption by facility over time. Clear charts make it obvious which facilities are the largest contributors to an individual's energy footprint (Figure 6.1).

The mobile phone is the world's most popular computing platform and it provides an excellent platform on which to develop a personal energy meter: the inbuilt sensors can provide much of the contextual information needed to personalise the results to an individual (Chapter 5) and by its very nature it is always carried and so can provide a well-timed 'nudge' about energy consumption at appropriate moments. This has been shown to be one of the most effective feedback strategies (Section 2.1). Furthermore, although input from many sensors and systems will be needed to arrive at each person's estimate, calculating the overall consumption on that person's mobile phone has important privacy-preserving properties, doing away with the requirement for any external system to be able to associate together too much personal information.

6.3 Situated subscription

In a world full of energy feeds, identifying the relevant ones could become a challenge. Searching central directories, or exchanging long addresses, is tedious, unfriendly and time-consuming. An added advantage of the mobile platform is that it provides a fast and effective mechanism for subscription: the example aggregator can subscribe to a feed just by scanning its URL encoded as a 2D barcode and physically attached to the object to whose energy consumption it refers (Figure 6.2(a)). One can therefore sign up

(a) the coffee machine's consumption (b) building energy consumption

Figure 6.2: QR codes allow users to subscribe to feeds

for feeds when one is close to them. New generations of phones will include near-field communication support, allowing another possibility for situated subscription. Each feed can include data for several different facilities which are offered to the user (Figure 6.3(a)), meaning a single code near the entrance to a building can refer a visitor to all the relevant energy information (Figure 6.2(b)). Additionally, facilities can be suggested based on location from a database of geotagged feeds. In these ways a user can build up 'on-the-fly' a list of subscriptions relevant to the main consumers of energy in his life (Figure 6.3(b)).

6.4 Feeds and case study

This section describes a case study involving three people who work in the William Gates Building to demonstrate what a personal energy meter might produce. Its information came from manual input as well as various sensor systems.

Figure 6.4 shows the result of this process for one person for a single week. On the left hand side are the raw sensor data inputs; the middle column shows the share of energy consumption of each type allocated to the individual, while the right-hand graph shows the output of this personal energy meter for the week and an energy stack comparable to Figure 3.1. The average daily consumption came to 135 kWh.

A number of feeds were built based on data from websites or devices described in Chapter 2, both to demonstrate how simple the process can be and to ensure that new users of a personal energy meter would be able to derive some benefit from its feedback without

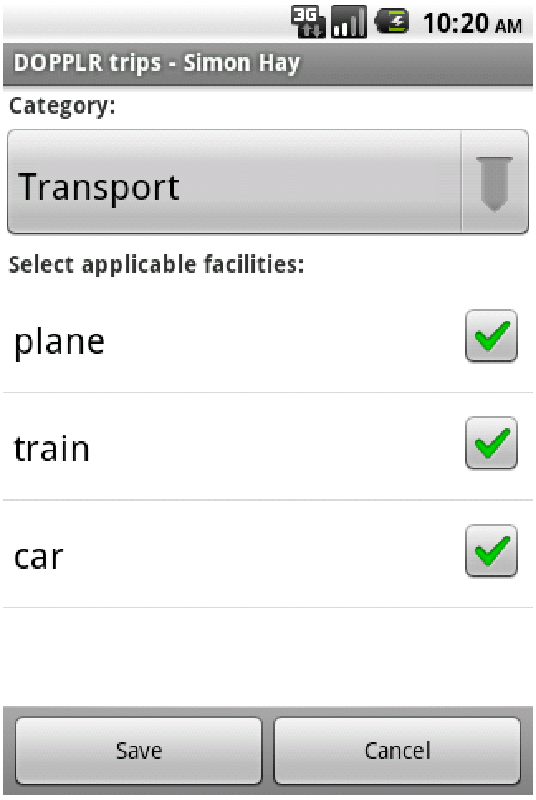

(a) each feed contains many facilities

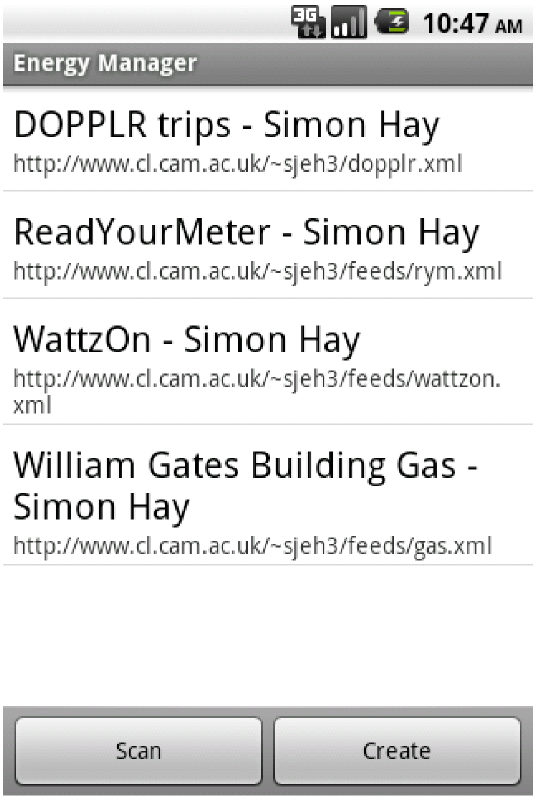

(b) feed subscriptions

Figure 6.3: Subscribing to feeds using Energy Meter

the need to deploy any additional sensors. This is in keeping with the principle of incremental sensing described in Chapter 1: where humans can themselves be used as sensors this can prove an effective starting point.

Feeds were also developed to estimate office energy consumption based on the techniques described in Chapters 3, 4 and 5.

6.4.1 Home energy consumption

ReadYourMeter[2] is a free website that allows users to record and graph their utility meter readings and so aims to help them understand their energy consumption, compare their data with others and see how much energy organisations use. Although over time many buildings will transition to smart meters that can report usage automatically, for now manually-entered readings remain vital to avoid excluding a large proportion of the population. The site was extended to include an energy feed for each user. Each feed contains a separate facility for each meter attached to the account.

Participants in the study recorded their home electricity, gas and water meter readings each day using ReadYourMeter.

[2] `http://readyourmeter.org/`

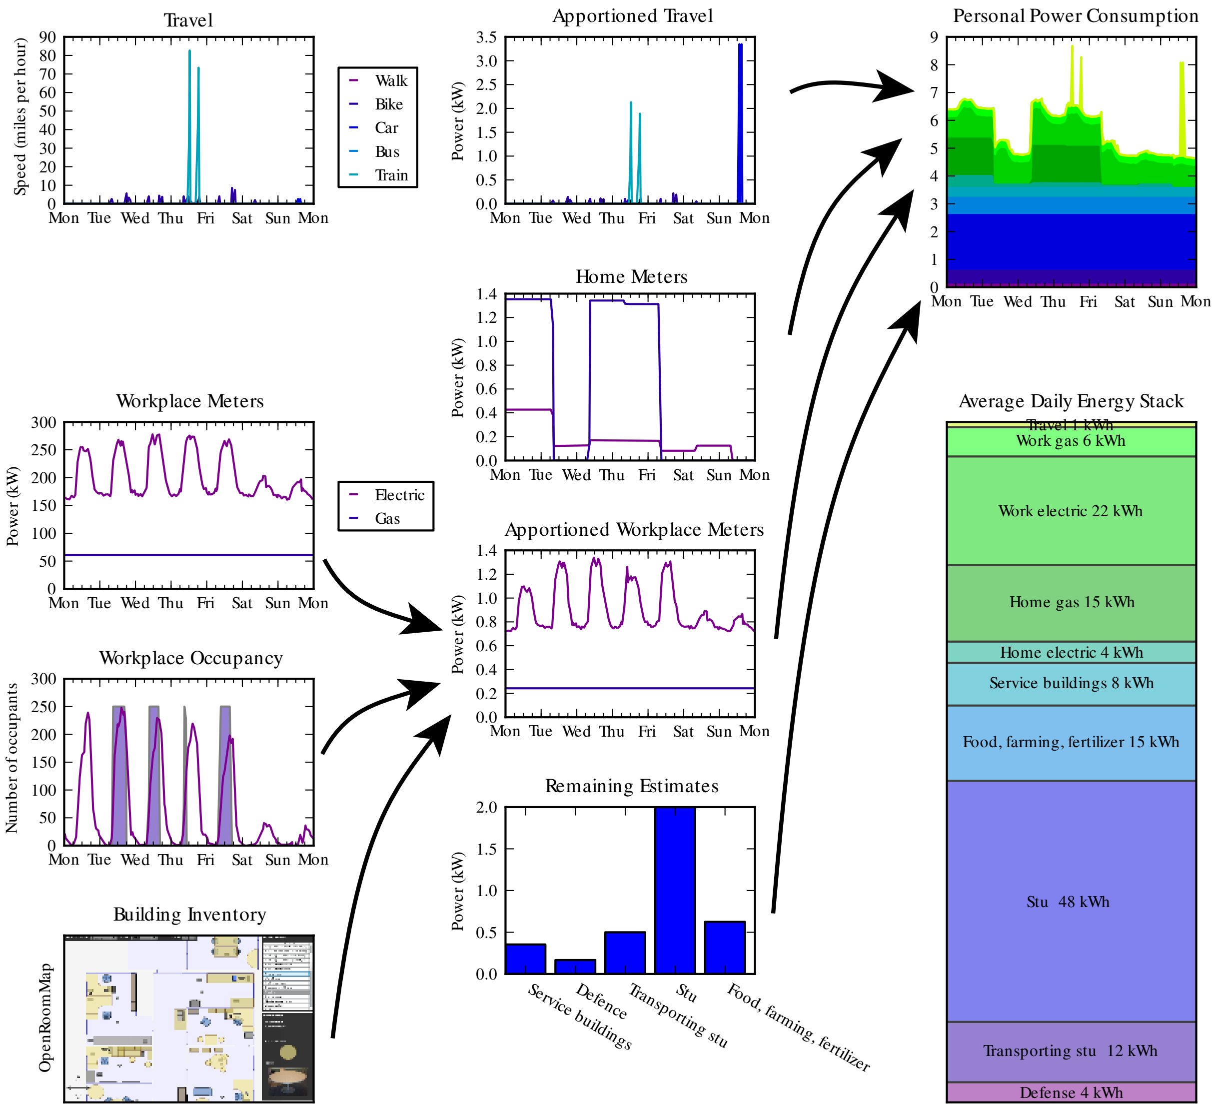

Figure 6.4: Inputs, processing and final result of personal energy metering

6.4.2 Transport

Energy consumption through transport accounts for around 35% of MacKay's estimate for the typical individual. Estimates can be made from location traces and knowledge of the mode of transport.

Participants in the study recorded GPS traces using Android-based mobile phones; these were verified using travel diaries and journey logs from the Sentient Van (Section 2.4.2). The mode of transport was identified manually by each participant and energy consumption for each journey estimated very simply using average figures suggested by MacKay; clearly, more sophisticated models could be used taking into account variables such as speed. An average UK car is assumed to do 33 miles per gallon, which corresponds to an energy consumption of 80 kWh per 100 km; MacKay also calculated that the (invariably full) Cambridge to London train has an equivalent energy consumption of 1.6 kWh per 100 passenger km. A simple static estimate of occupancy was used to estimate the energy consumption of trains and buses, but the principles learned from studying the energy consumption of large buildings (Section 3.5) could be applied to investigate more accurate apportionment of the energy costs of public transport. In the same way that dividing the energy costs of a building amongst its occupants means one person working for a paper deadline is allocated all the power used for heating and servers, it is unreasonable just to divide the energy cost of a bus evenly amongst all those on it: a late-night bus might only have a few passengers but if it did not exist many more people might have driven into work rather than worry they might not get home if a meeting ran late.

Even cycling has an energy cost, estimated by MacKay as around 1 kWh per 100 km — the energy expenditure of the person himself. This value is used for the traces in this study. There are a number of parallels between our own energy expenditure and our use of external energy: in general, we use energy-consuming machines to reduce the amount of work we have to do ourselves and augment our capabilities, and so it makes sense to consider the two concepts on the same scale. Furthermore, the same sensors can often be used to derive both sorts of information.

It is informative to look at how badly the overall estimate is degraded if such fine-grained sensor data is not available. For example, there may only be a journey log rather than a full location trace, or people might manually enter odometer or fuel readings at sporadic intervals in the same way they do electricity meter readings. GPS traces of journeys by bicycle were examined and the measured distance compared with an estimate from the Cycle Streets journey planner.[3] The discrepancy ranged from almost nothing to approximately 10%. This compares favourably with the uncertainty involved in converting distance travelled to energy costs in the first place. The error was higher for the train journeys since distances by rail are not generally published; if bus and train companies share this information as part of their timetables it would be possible to make more accurate estimates.

6.4.2.1 Dopplr

Dopplr's API provides dates, destinations and modes of transport for trips users have entered into the site (see Section 2.4.2).[4] The Dopplr feed for a user contains a facility for each mode of transport, updated whenever a trip is taken; the energy cost of the journey is estimated by calculating the great circle distance and multiplying it by a factor depending on the vehicle [136]. This will be inaccurate for short journeys, but a much better approximation for long-haul flights; this seems reasonable since most people tend to use social sites like Dopplr only for longer, more significant trips. Clearly, more intelligent mechanisms could be applied to estimate both distance and energy cost more accurately (such as route finding algorithms; see Section 3.2), but this simple estimate is sufficient to demonstrate potential.

None of the participants in this study travelled by air in the week in question, so this category does not appear in the results presented.

6.4.3 Office energy consumption

Feeds were also built for the energy consumption of the William Gates Building based on the techniques presented in this dissertation.

[3] `http://www.cyclestreets.net/`
[4] `http://www.dopplr.com/`

The electricity meter of the office building, in common with those of many large buildings, logs half-hourly measurements of the total energy consumed (Section 3.5). Based on the results of the apportionment study described in Section 3.5, a *personal load* policy was used for the allocation of the workplace's electricity consumption. This policy meets the apportionment principles by assigning personally attributable energy load for to each occupant in the building and then dividing the remainder evenly amongst all occupiers (currently present in the building or otherwise).

The policy was implemented with the initial estimate that all users of the building have the same personal load. This starting point gives the opportunity to refine incrementally the personal load estimate as resources permit. For example, a network monitor and inventory of computer systems could be applied to assign the actual energy cost of a workstation (perhaps left on overnight) to its user.

Implementing this policy requires both knowledge of the current occupants and their personal load estimate. In the William Gates Building where this study was focussed, logs from the access control system were therefore used to estimate the number of people inside at any given time, in the manner described in 3.5.3.1, since most of these people were not participating in the study. In keeping with the principle of making sensing optional to improve accuracy, rather than mandatory, feeds were created for the example working patterns described in Section 3.5.1 using a standard estimate of personal load. Personalised feeds are also available which make use of location data, such as that provided by the system described in Section 5.1.1, and more accurate estimates of personal load based on OpenRoomMap ownership information (see Section 2.5.2). Finally, a similar set of feeds was created using the data from the model described in Section 4.2 rather than direct measurements; this has the twin advantages of not requiring any sensing and providing an approximate breakdown by function.

The gas consumption of the building is for heating and hot water from which all occupants benefit. In particular the thermal inertia of the building means that heating even benefits occupants who arrive after it is switched off. A simple static apportionment policy was therefore applied to the building's gas consumption to allocate a fixed proportion to all potential occupants of the building. This has the advantage that only a single feed is required; participants can subscribe from a printed tag on a noticeboard.

6.4.4 Remaining fixed estimates

WattzOn[5] allows users to estimate their total energy footprint by answering a series of questionnaires with the stated goal of educating users about energy efficiency and conservation. It also features a embodied energy database[6] containing details of the footprints of a significant number of consumer goods. Users can select the items they own to have their costs added to their profiles. All this data is also exposed through an API, making it straightforward to write a feed providing estimates of consumption for housing, food, commuting, flying, government and 'stuff'. The government, food and stuff estimates were used in the case study to account for items not within the scope of the dynamic monitoring.

[5] `http://www.wattzon.com/`
[6] `http://www.wattzon.com/stuff`

6.4.5 Additional feeds

6.4.5.1 Plogg

At the opposite end of the spectrum from meter readings in terms of both granularity and timing, the architecture is also appropriate for monitoring the consumption of individual appliances. A feed was created to interface with the smart plugs produced by Plogg (see Section 2.3.1.3) and provide readings once a minute. Future devices might have this functionality built in, or might report their readings back to a separate web service which provides a feed in similar manner to Google PowerMeter.[7]

6.5 Summary

This chapter has presented a lightweight architecture for federating consumption data from multiple disparate sources and demonstrated a proof-of-concept in the form of an aggregator for mobile phones and several feeds based on the technologies developed in previous chapters.

[7] `http://www.google.com/powermeter/`

Chapter 7

Conclusions

Contents

Overview

This chapter revisits the research questions posed in Section 1.4, outlining possible avenues for future research and summarising the main contributions of this dissertation.

7.1 Research contributions

This dissertation has made the following contributions:

1. A description of the novel concept, benefits and principles of a personal energy meter (Chapter 1) and a literature review drawing together disparate threads of existing work to show the basis for a personal energy meter (Chapter 2).

2. A presentation of the case for and concepts of apportionment, including a taxonomy of resources and the data requirements to handle them; an evaluation of a range of strategies in a case study and elaboration of the overriding principles that are generally applicable (Section 3.5).

3. A lightweight, scalable approach to using only crowd-sourced inventories and device profiles to estimate building energy consumption (Section 4.2).

4. A mechanism for profiling devices to determine the energy costs of specific activities, particularly applicable to shared programmable devices (Section 4.5).

5. A description and evaluation of the new concept of inquiry-free Bluetooth tracking that has the potential to offer indoor location information useful for personal energy metering with significantly less infrastructure and calibration than other systems (Chapter 5).

6. A suitable architecture for a personal energy meter on a global scale and a demonstration using a mobile phone application to aggregate energy feeds based on the case studies and technologies developed earlier (Chapter 6).

7.2 Research questions revisited

To what extent can technology be used to apportion personal energy costs? Chapter 3 showed that consumption data and context can be combined to allocate the energy costs of shared resources to individuals.

Can energy consumption be inferred without continuous metering? Chapter 4 demonstrated how a building's energy usage can be modelled sufficiently accurately using only limited sensing, crowd-sourced inventories and device profiles.

Can context be monitored with minimal additional infrastructure? Chapter 5 presented novel Bluetooth-based indoor location systems that can provide the room-level context information using existing PCs and without extensive calibration.

What should be the software architecture of a personal energy meter? Chapter 6 suggested energy feeds to federate consumption data from disparate sensor systems and showed a model system in which a mobile phone aggregates information from all the systems discussed in this dissertation.

7.3 Further work

There are many potential ways to improve the building energy model described in Section 4.2. Useful examples might be to modulate building lighting with reference to natural light levels using more detailed weather data and information from OpenRoomMap about the positions of windows in offices. OpenRoomMap data could also provide more assistance in estimating building parameters by providing estimates of building surface area and the relative ratios of walls, windows and roofing. It would also be useful to evaluate the model's performance in other buildings. The simple physics engine was designed to show how little data is necessary to produce valuable results; it could be replaced with a more sophisticated model as used in other packages provided little additional input is required.

The tracking system described in Section 5.1.1 has very little in the way of deployment costs, allowing the construction of large testbeds. It would be useful to evaluate different tracking algorithms by deploying such a testbed and encouraging its use.

Considering the model personal energy meter as a whole, it would of course be beneficial to deploy the system more widely and obtain feedback from a wider group of users. It would be particularly valuable to encourage third parties to deploy energy feeds or create and evaluate their own aggregators with different visualisations or feedback mechanisms. There are significant HCI questions in how best to present data to be persuasive and how to help users share their results with others; for example, it would be valuable to calculate and display estimates of error for each data source in the aggregator.

7.4 Summary

A personal energy meter that provides live information on consumption apportioned to individuals represents a very significant step forwards from the current common situation of a static, approximate and time-consuming audit of a building or organisation. It is dependent on developments in a number of computing technologies—in particular, sensors and sensor networks to provide data both on usage and on interactions and a common world model to allow information to be collected wherever the user might be. It promises to provide important insights and incentives to help us each control our own footprint.

Apportionment is important and the correct choice of policy merits careful consideration. Different policies have significant effects on the total energy allocated to individuals. Nevertheless, all apportionment policies should exhibit *completeness*, *accountability* and *social efficiency*. *Personal load* provides the best opportunity to personalise results and improve accuracy incrementally and offers valuable incentives for users to reduce their consumption. The two crucial ingredients for apportionment are metering and context.

Low fidelity sensing, or in many cases just prior knowledge and public data sources, could still result in overall estimates with acceptable error; the ability to tolerate incremental sensor deployment is necessary to encourage widespread adoption. Estimates of building energy consumption can be formed from device profiles and inventories. Minimising the effort involved in initial data collection is important and the inventory data can be crowd-sourced from building users. Fine-grained profiles of programmable devices can be obtained using the framework presented which embeds synchronisation information in the measurement trace itself, making the entire process automated and repeatable.

Location is the best form of context, but few indoor location systems are suitable for widespread deployment outside research environments due to the extensive bespoke infrastructure that must be installed and surveyed. This is costly in terms of both money and time, and impractical in most buildings. Bluetooth-based systems are an attractive alternative as they have low power requirements and almost everyone already carries a mobile phone and has a computer on his desk; the use of low-level Bluetooth connections to track mobile devices within a field of fixed base stations has been shown to be a viable technique for the construction of a low-cost, low-infrastructure location system that can be deployed globally.

Data from disparate sensor systems must be brought together to build a complete picture of an individual's energy consumption. This can be achieved using a framework based on the Atom Syndication Format and aggregator applications; this has privacy and simplicity advantages over centralised systems. The mobile phone is an ideal platform for a model aggregator since its inbuilt sensors can provide much of the contextual information needed and billions of people carry one all the time.

The sensing techniques and algorithms presented here have been demonstrated with power but also be applied to apportion ecological footprints or carbon externalities; they form the basis of a truly general personal energy meter.

Appendix A

Energy consumption for wireless networking

Contents

A.1 Mobile phone consumption

A study of the power consumed by sending data over a wireless network from a mobile phone provides a demonstration and evaluation of the measurement platform described in Section 4.5 and shows the level of detailed understanding of consumption that it makes possible. Since mobile phones are generally optimised for power consumption already, techniques applicable to them are likely also to apply elsewhere.

Pervasive computing is a vision of communicating devices and so understanding energy costs of this communication is also of great importance to application developers. The interaction between different layers in the hardware and software stack creates considerable differences in energy consumption, which provides significant motivation for measurement frameworks such as this. The mobile phone is a particularly appealing platform for pervasive computing applications, but as with all battery-powered devices, controlling and managing power consumption is an issue. These applications have a number of notable characteristics with respect to power consumption. Firstly, many context-sensitive applications will run continuously in the background collecting sensor information or waiting

Some of the contributions presented in this appendix have also appeared in separate publications [175, 176].

for a trigger condition and so even a small power requirement has the potential to impact the device more heavily than other power hungry but short-lived programs. Secondly, many applications rely on the participation of a large group of users to be useful. Perhaps one of the most compelling cases is the concept of participatory sensing where a large number of volunteers can use their smart phones to gather sensor data in the background [24]. This can subsequently be used for all sorts of purposes, such as to build up a picture of environmental factors [147] or generate collaborative maps [38]. For these applications to succeed the cost of participation must be small compared with the direct benefit gained [141] and so the power impact of running the application must be minimised. Finally, these applications often operate in varying conditions but with flexibility about how a particular task is performed—for example, sensor readings can be reported immediately or stored up for later transmission. Applications should choose the best approach from the various options available and dynamically integrate with the overall usage of the device.

A.2 Network traffic monitoring

It is useful to observe the network traffic alongside the power trace of the mobile device in order to analyse the costs of different methods of sending and receiving data. The PC recording the power trace was connected to a wireless access point and configured to run a DHCP server to emulate a typical network the phone might join. The Power Server application then called `libpcap`[1] to record all packets seen on that interface.

The framework combines the synchronisation information embedded in the power trace with the timing log from the phone and the network traffic information in order to generated annotated graphs of power consumption. An example of this output is shown in Figure A.1 and discussed in the remainder of this appendix.

A.3 Connecting to the network

Figure A.1 shows part of a measured energy trace when connecting to a WiFi network. The top trace shows the cost of obtaining an IP address using DHCP, annotated using the method described with each IP packet sent and received. The bottom trace shows the same operation but using static addressing. Without the packet labels this trace would be relatively hard to interpret, but the aligned annotations show clearly the costs of each aspect of the connection process. Note that the repeated DNS requests come from the operating system itself and are for a Google server; Android attempts to make contact at regular intervals, and these communications also show up in other test runs, distorting the results. The ability to identify and account for these occurrences is another advantage of the annotation system. The actions taken by the phone when connecting to the network are prescribed by the various Internet RFCs. For example, the ARP probe packets are designed to discover if there is another host on the network already using the phone's desired IP address [168]. The number of probe packets and the delay between them account for a significant fraction of the connection time (and energy).

[1] `http://www.tcpdump.org/`

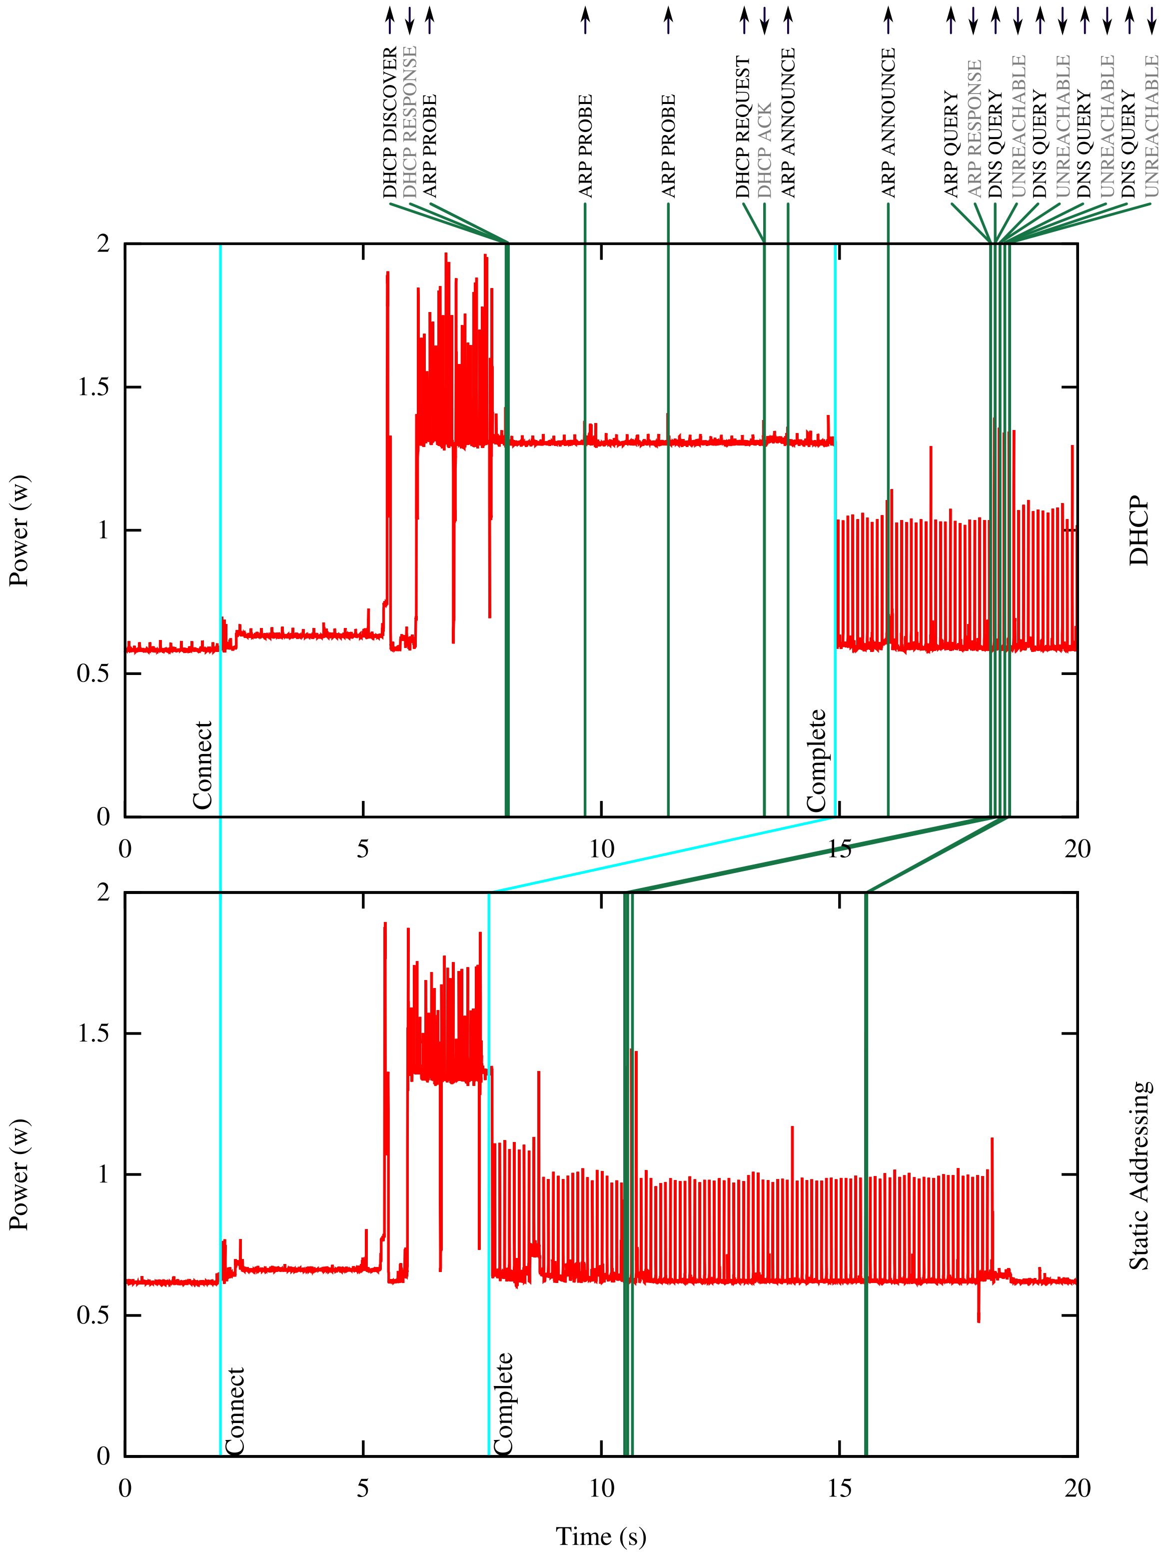

Figure A.1: Energy trace of connecting a G1 handset to a WiFi network [176]

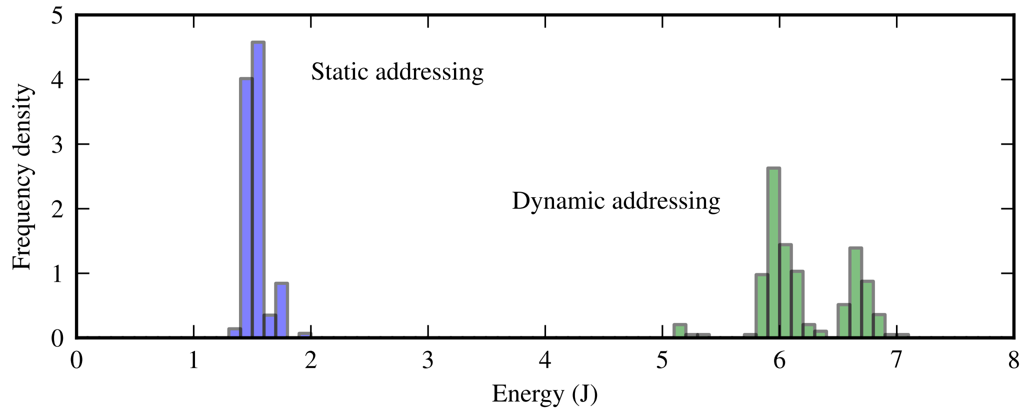

Figure A.2: Energy consumed by a G1 handset connecting to the wireless network [176]

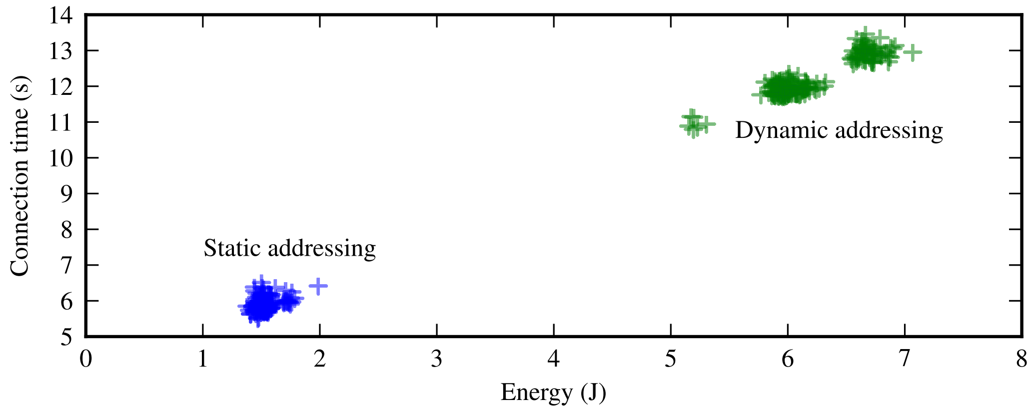

Figure A.3: Energy against time-taken to connect to a wireless network [176]

A considerable energy saving is available from the use of static addressing. The histogram in Figure A.2 shows a summary of the energy consumed in 200 trials of each of these two techniques. The most common cost for a dynamic connection is around 6 J whereas a static connection commonly consumes only 1.5 J. It is also notable that the energy consumed by dynamic addressing seems to be partitioned into a number of separate distributions. This is in fact due to discontinuities in the time taken for a connection to complete as shown in Figure A.3. The cause of this phenomenon is not yet clear. One possible explanation is that of a timeout in a polling loop in the operating system or the power control hardware in the wireless chip. The clusters seen in Figure A.3 are at approximately 1 s separation—one could imagine this as an appealing number to a developer who needs to select a value for a timeout or sleep period.

One means to improve the energy efficiency of connection whilst maintaining the flexibility of dynamic addressing is to eliminate the ARP probe stage from the process. This is permitted by the RFC in specific situations [168]. In fact this optimisation has been applied in later versions of the Android operating system. Figure A.4 shows the connection traces for a Google Nexus handset connecting to a wireless network.

This saving is clearly evident in the histogram of connection cost for the Nexus handset (Figure A.5). Note also that the use of static addressing on this handset continues to

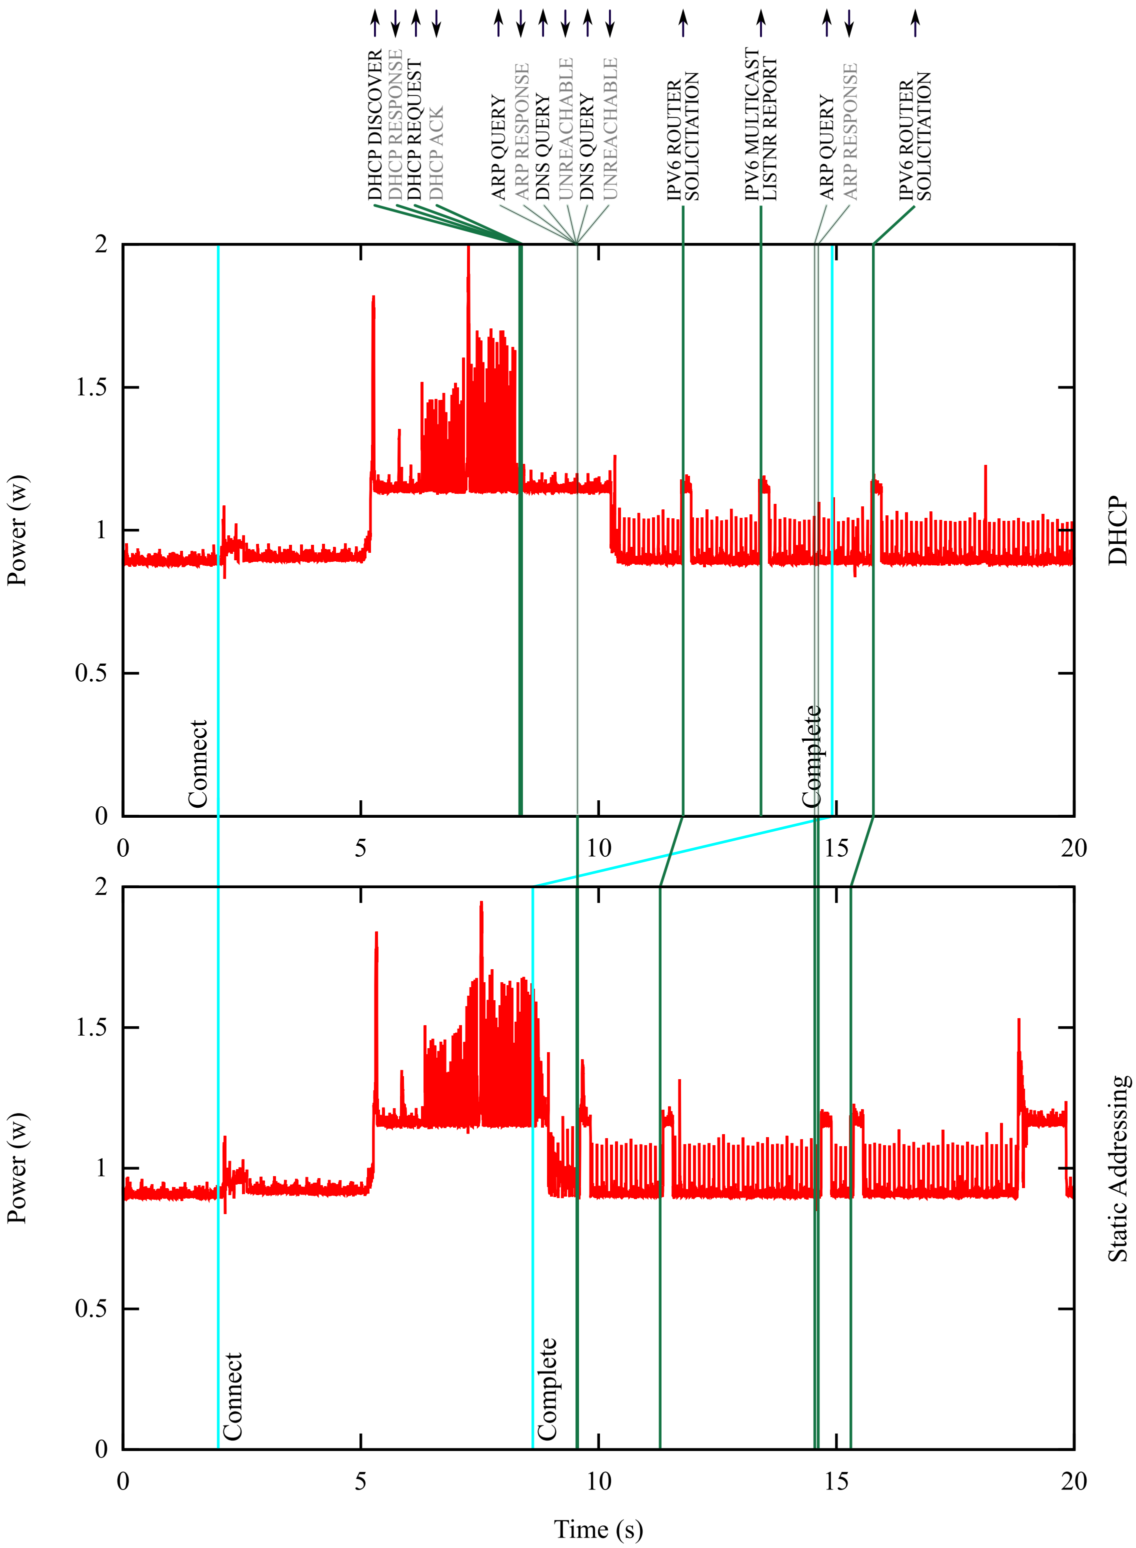

Figure A.4: Energy trace of connecting a Nexus One handset to a WiFi network [176]

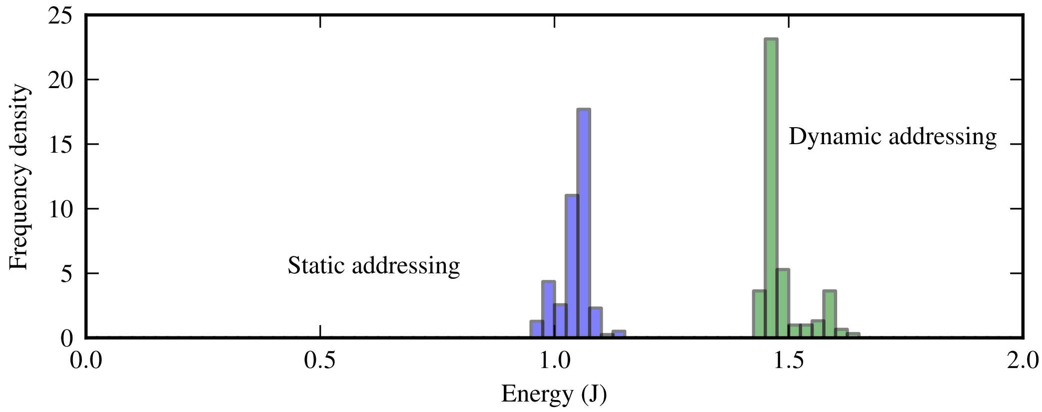

Figure A.5: Energy consumed by a Nexus One handset connecting to the wireless network [176]

present an energy saving albeit highly reduced.

A.4 Energy saving in context

Owners of G1 handsets might consider switching to static addressing. The typical saving in this case would be around 5 J. This is approximately equivalent to:

5 seconds of talk time The average power consumption when on a call seems to be around 1 W in an office in Cambridge (though this may vary with distance to cell towers).

500 seconds (8 minutes) of standby time The average power consumption with the phone in standby is around 0.01 W (Figure 5.16).

200 seconds (3.5 minutes) of idle WiFi connection An idle WiFi connection adds around 0.024 W to the phone's consumption (Section A.5).

Alternatively, a device which makes a connection every 10 minutes (for example polling for new email) makes around 144 connections a day. Assuming a nightly charging strategy with a typical battery of 1,400 mAh at 3.7 V the saving corresponds to around 4% of the total battery life of the handset.

A.5 Idle power

Figure A.6 shows excerpts of the energy trace of a handset when connected to only the 3G, 2G or WiFi networks. The trace shows only the consumption of the wireless networking hardware because the baseline power consumption of the phone (CPU, backlight etc.) have been subtracted from the trace. In this case WiFi actually has the lowest idle power cost, followed by 3G and then 2G. Although the spikes are more frequent (every 100 ms,

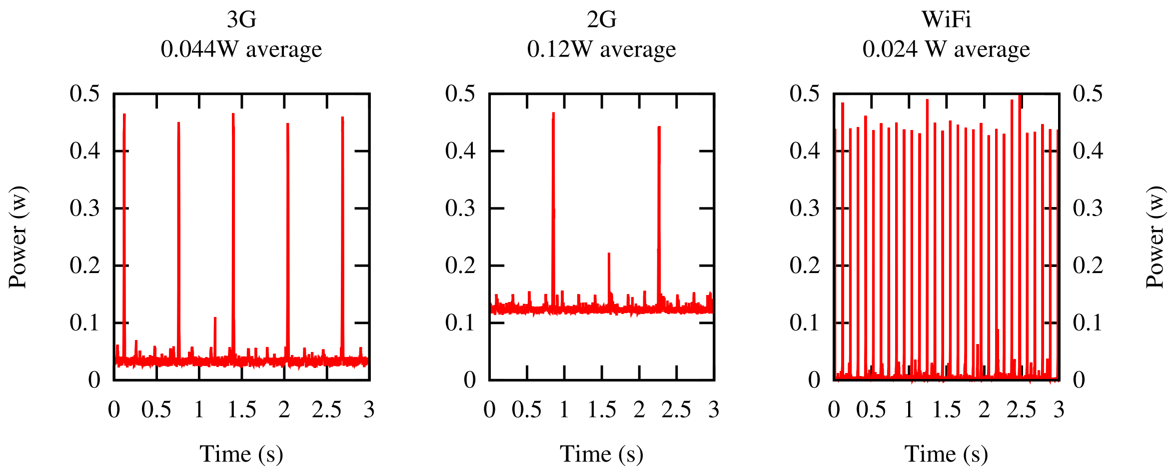

Figure A.6: Additional power consumption incurred when connected to 3G, 2G and WiFi networks [176]

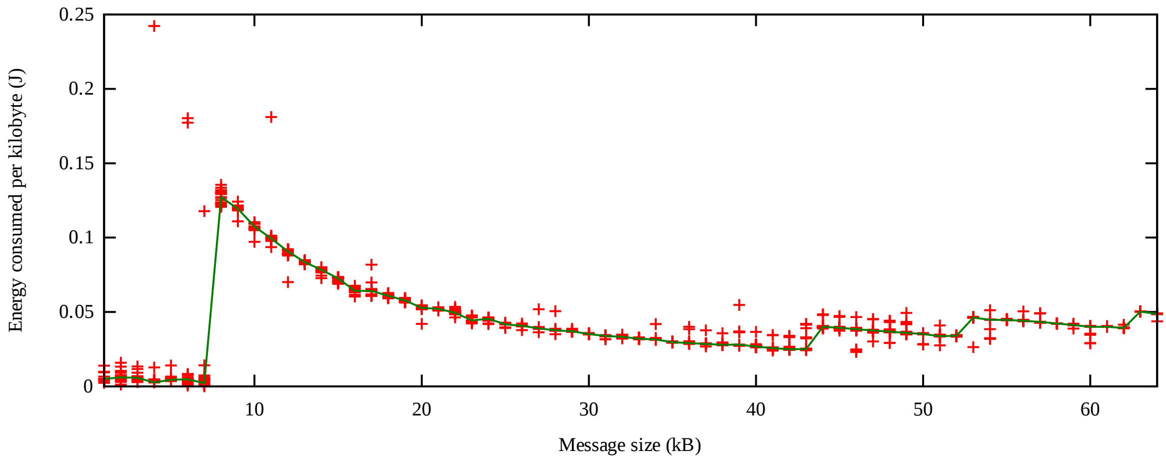

Figure A.7: Variation in G1 energy cost per unit data with total message size [176]

corresponding to receiving base station beacons), the base power is lower than maintaining a connection to the cellular network.

It is not yet possible to draw concrete conclusions from these measurements since there are several potential factors which have not been investigated. The locations of the various base stations will have an effect on the power consumed by the radio and the building itself will have different attenuation properties at the different radio frequencies involved.

However, these measurements do demonstrate that one cannot always assume that one particular networking technology will have the lowest power consumption. For example, 2G networking is provided as an option in the phone's interface to reduce power consumption but in this particular case it is the highest power option. The power measurement framework provides sufficiently detailed information to allow assumptions such as this to be questioned and investigated.

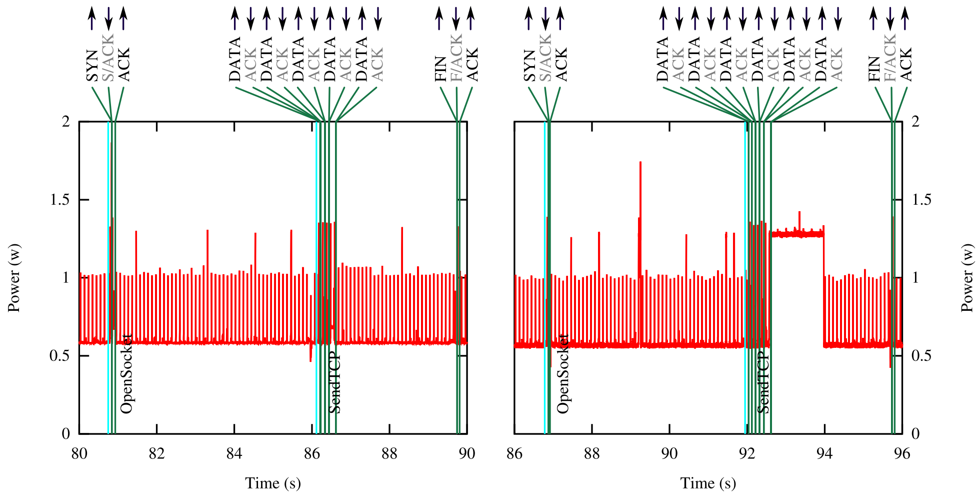

Figure A.8: Extracts from the G1 energy traces of sending 7 KB (left) and 8 KB (right) of data over WiFi [176]

A.6 Data transmission

One might intuitively expect that the energy cost of sending a byte of data reduces, or at worst remains constant, as the total amount of data sent increases. However, this is not the case. Figure A.7 shows the number of joules required to transmit each kilobyte of data for increasing total message size using the G1 handset. The baseline calculation functionality in the framework was used to remove the residual costs of running the phone and so these numbers are the actual amount of additional energy required to send the data. The graph shows the result of 10 test repetitions run at each 1 KB interval.

Part of the reason for the noise in this data is that other processes on the phone are also using the network. In this test case they are attempting DNS look ups of particular Google servers. The evidence of this activity was visible in the packet trace collected by the Power Server.

There is a clear jump in the cost per byte for 7 KB of data compared with 8 KB. Figure A.8 shows these two instances in more detail. When sending 8 KB or more of data (Figure A.8(right)) there is a considerable period of high energy consumption after the last packet has been sent which is not present when sending a 7 KB message (Figure A.8(left)).

It is not clear why this is occurring and there is no explanation in any of the relevant networking standards. Use of a separate wireless card and the Wireshark[2] packet sniffer demonstrated that there is no activity on the wireless network for this period. However, regardless of whether this is due to crossing some power management threshold or simply a bug in the wireless firmware, it has a significant impact on the energy requirements of sending a message. A pervasive sensing application on one of these devices would minimise power by batching a data update into chunks of around 7 KB but incur a significantly larger cost by batching to 8 KB. These results are also almost identical when using the HTC Magic handset and the HTC Hero handsets.

The network proximity of the test server process to the phone means that the results are for a TCP connection with very low Round Trip Time (RTT). This fact combined

[2] http://www.wireshark.org/

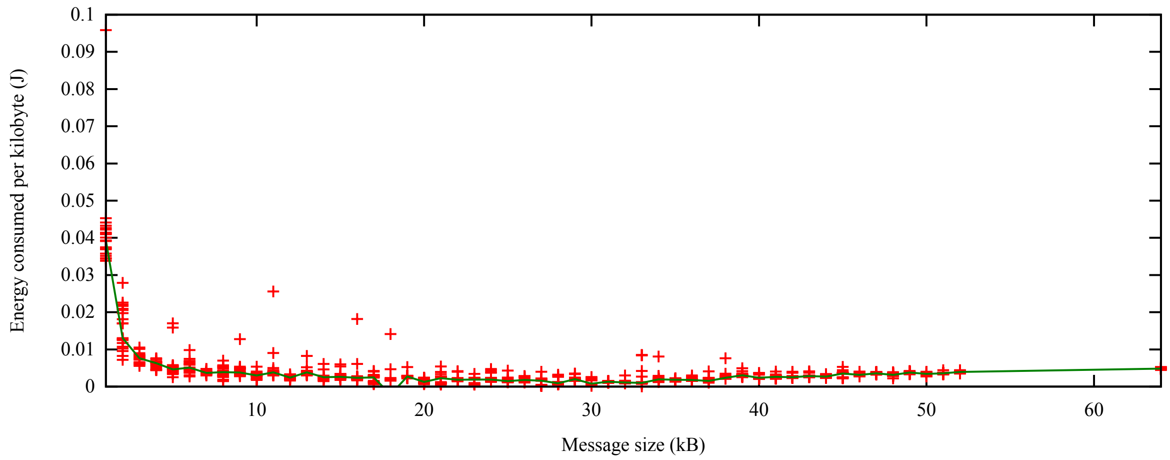

Figure A.9: Variation in N1 energy cost per unit data with total message size [176]

with the high packet processing speed of the test server means that the phone rarely manages to send a second packet before the acknowledgement of the first packet arrives. This assertion was validated by running a set of automated tests with varying connection latency. High latency connections (of the order of 100 ms RTT) show conventional TCP slow start behaviour.

Figure A.9 shows the variation in energy cost for the Nexus handset. The trend shown in this graph conforms much more to expectations in which longer messages are more efficient than shorter ones. When compared with the G1 results (Figure A.7) it is clear that the best case cost for both handsets is around 0.005 J per kilobyte of data, whereas the worst case for the G1 is approximately 0.13 J/KB and the worst case for the N1 is approximately 0.04 J/KB. These extremes occur in different places on the graph for the two handsets. In fact a message size of 8 KB is close to the best efficiency on the N1 and the worst on the G1.

A.7 Send buffer size

As a final example the impact of the size of the send buffer used by the application developer is assessed. Android applications are written in the Java programming language and network data is sent by getting the `OutputStream` object associated with a `Socket` instance. Data is then sent over the network by calling the write method on the socket and passing an array of bytes to send.

Nagle's algorithm is used in TCP to aggregate small data chunks into a single larger packet [148]. However, this occurs only when there is unacknowledged data in transit. Given the small RTT of the test setup this is rarely the case and so the byte array is sent immediately to the client socket without waiting for further data. The size of this array therefore causes significant changes in energy costs. Figure A.10 shows how the energy cost per kilobyte varies with changes in the size of the byte array passed from the application for a message of 1 KB and a message of 32 KB. Note that the buffer size (on the horizontal axis) is shown on a log scale. For both of these messages the choice of buffer size can cause a tenfold difference in the energy cost!

One would expect the energy cost per KB for the 32 KB message to be lower, since the

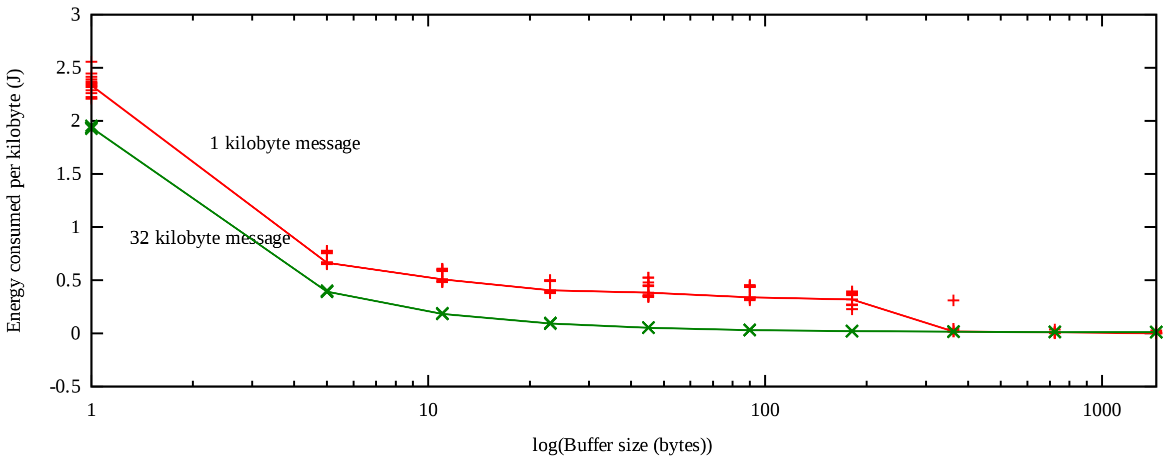

Figure A.10: Variation in energy cost per unit data with buffer size [176]

fixed costs are amortised over more data. However, the 1 KB line collapses from its flat trend to the same level as the 32 KB line beyond a certain buffer size. This is because of the 7 KB/8 KB barrier described in the previous section: the change in cost is due to the number of frames sent over the Ethernet regardless of how much data is in them, and with a large enough send buffer there are sufficiently few frames that the 1 KB message does not cross this barrier.

A send buffer of 1,448 bytes maximally fills the payload of a TCP packet when sent over Ethernet. The best (and most consistent) performance results are when using a send buffer of this size. If the buffer is smaller than this then transmission efficiency is lower because packets are sent only partially filled. If the buffer is larger than this then the operating system must fragment data over multiple packets—this operation also incurs an overhead and reduces transmit performance.

Appendix B

Circuit diagrams

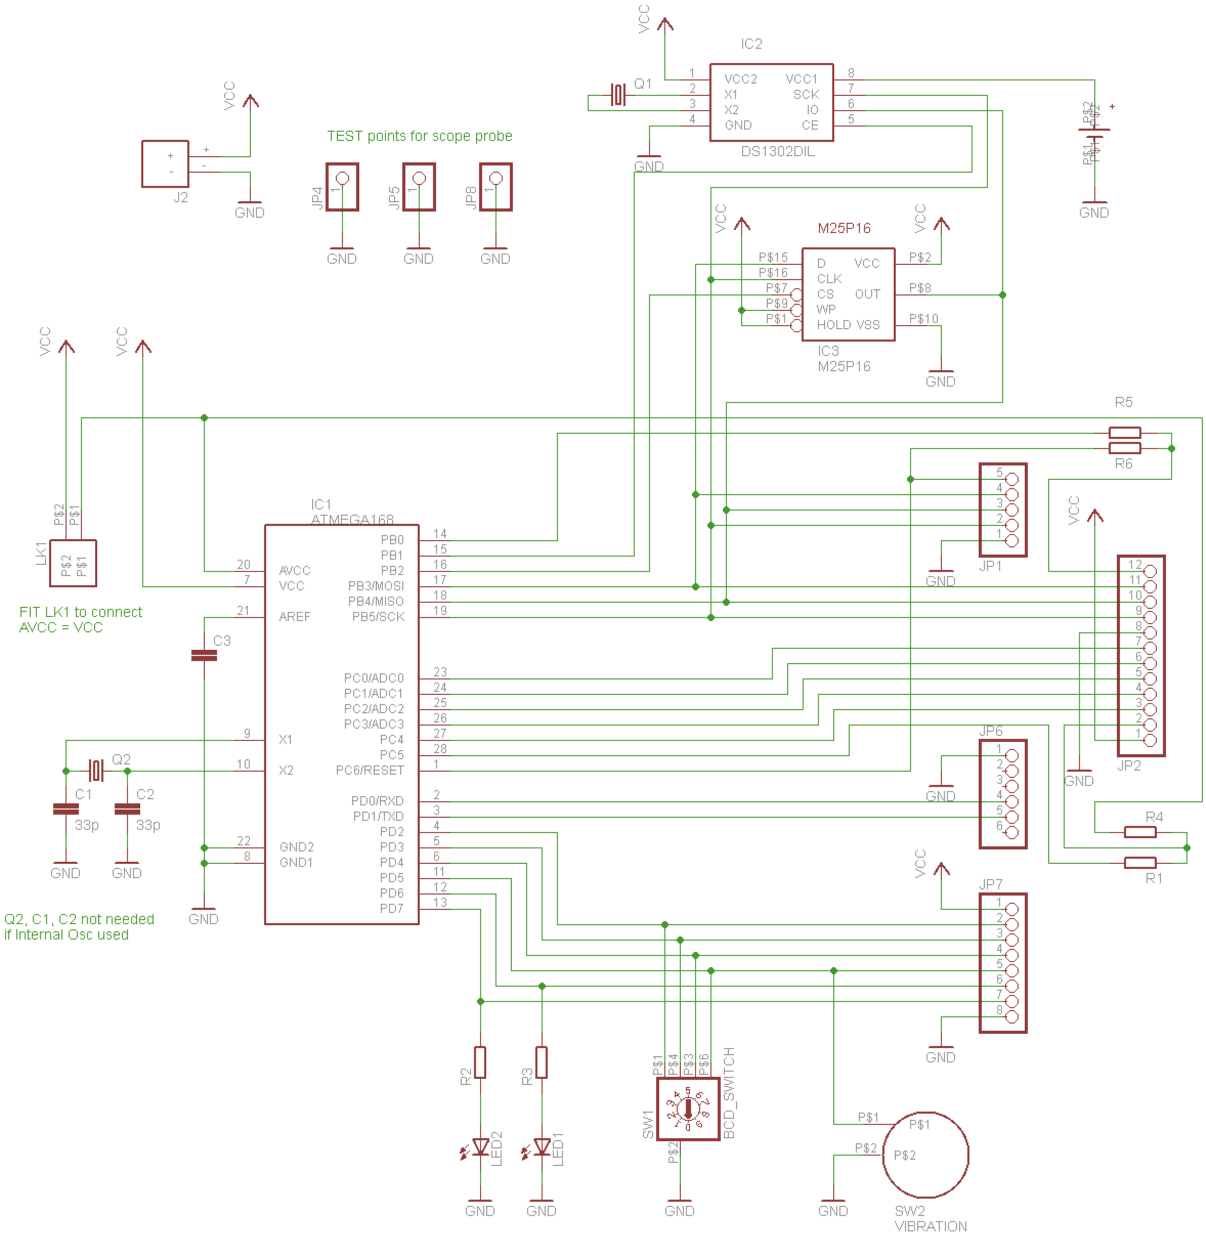

Figure B.1: Schematic circuit diagram for light sensor node described in Section 4.1.1

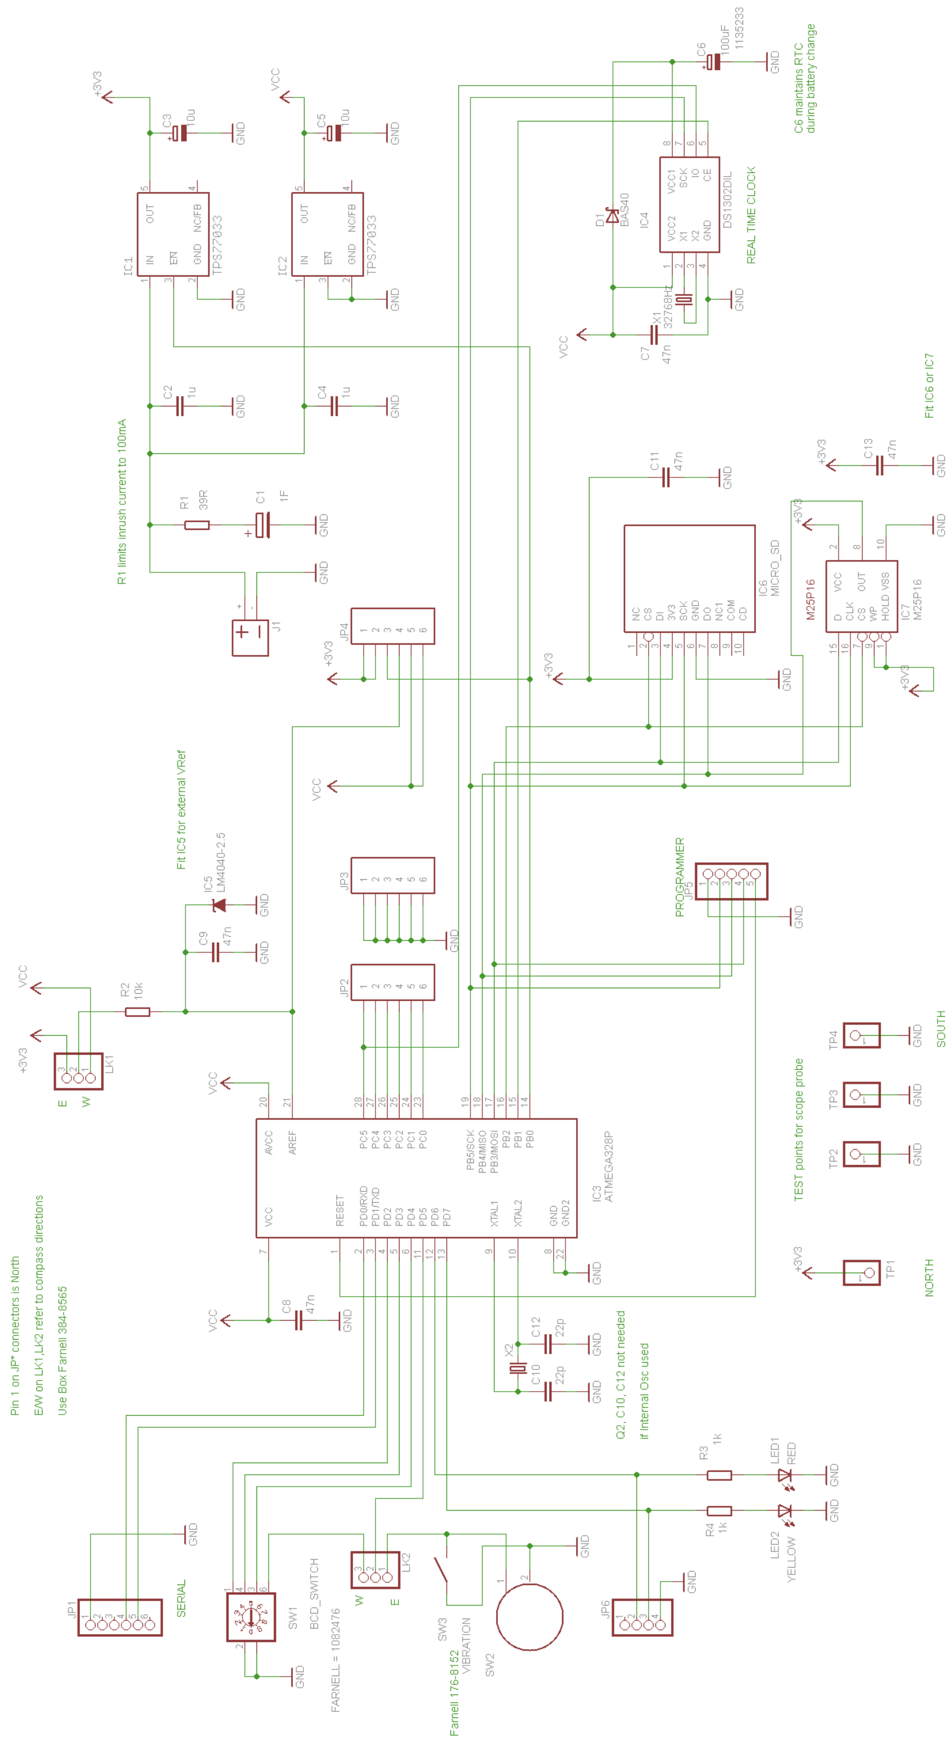

Figure B.2: Schematic circuit diagram of measurement hardware described in Section 4.4

References

[1] W Abrahamse, L Steg, C Viek, and T Rothengatter. A review of intervention studies aimed at household energy conservation. *Journal of Environmental Psychology*, 25(3):273—291, 2005.

Cited on page: 32.

[2] N Adams, R Gold, B. N Schilit, M Tso, and R Want. An infrared network for mobile computers. *MLCS '93: Proceedings of the 1993 USENIX Symposium on Mobile and Location-independent Computing*, pages 41—52, 1993. USENIX, Berkeley, CA, USA.

Cited on page: 58.

[3] M Addlesee, R Curwen, S Hodges, J. F Newman, P Steggles, A Ward, and A Hopper. Implementing a sentient computing system. *Computer*, 34(8):50—56, Sep 2001.

Cited on pages: 59 and 145.

[4] M Addlesee, A Jones, F Livesey, and F Samaria. The ORL active floor. *IEEE Personal Communications*, 4(5):35—41, Oct 1997.

Cited on page: 56.

[5] Y Agarwal, S Hodges, R Chandra, J Scott, P Bahl, and R. K Gupta. Somniloquy: augmenting network interfaces to reduce PC energy usage. *NSDI '09: Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation*, pages 365—380, 2009. USENIX, Berkeley, CA, USA.

Cited on page: 39.

[6] Y Agarwal, S Savage, and R. K Gupta. Sleepserver: A software-only approach for reducing the energy consumption of PCs within enterprise environments. *USENIX '10: Proceedings of the USENIX Annual Technical Conference*, page 22, 2010. USENIX, Berkeley, CA, USA.

Cited on page: 39.

[7] Y Agarwal, T Weng, and R. K Gupta. The energy dashboard: Improving the visibility of energy consumption at a campus-wide scale. *BuildSys '09: Proceedings of the 1st ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, in conjunction with SenSys '09*, pages 55—60, 2009. ACM, New York, NY, USA.

Cited on page: 39.

[8] G Anastasi, R Bandelloni, M Conti, F Delmastro, E Gregori, and G Mainetto. Experimenting an indoor Bluetooth-based positioning service. *ICDCSW '03: Proceedings of the 23rd IEEE International Conference on Distributed Computing Systems Workshops*, page 480, 2003. IEEE, Washington, DC, USA.

Cited on page: 72.

[9] R. T Azuma. A survey of augmented reality. *Presence: Teleoperators and Virtual Environments*, 6(4):355–385, 1997.

Cited on page: 38.

[10] P Bahl and V Padmanabhan. RADAR: an in-building RF-based user location and tracking system. *INFOCOM '00: Proceedings of the 19th Annual Joint Conference of the IEEE Computer and Communications Societies*, pages 775—784, 2000. IEEE, Washington, DC, USA.

Cited on page: 68.

[11] Y Bai and C Hung. Remote power on/off control and current measurement for home electric outlets based on a low-power embedded board and ZigBee communication. *ISCE '08: Proceedings of the 2008 IEEE International Symposium on Consumer Electronics*, 2008. IEEE, Washington, DC, USA.

Cited on page: 41.

[12] M Bang, C Torstensson, and C Katzeff. The PowerHouse: A persuasive computer game designed to raise awareness of domestic energy consumption. *Pervasive '06: Proceedings of the 4th International Conference on Pervasive Computing. Lecture Notes in Computer Science, 2006, Volume 3968*, pages 123–132, 2006. Springer, Heidelberg, Germany.

Cited on page: 37.

[13] M Bargh and R de Groote. Indoor localization based on response rate of Bluetooth inquiries. *MELT '08: Proceedings of the 1st ACM International Workshop on Mobile Entity Localization and Tracking in GPS-less Environments*, pages 49—54, 2008. ACM, New York, NY, USA.

Cited on pages: 72 and 74.

[14] A Barry, B Fisher, and M Chang. A long-duration study of user-trained 802.11 localization. *MELT '09: Proceedings of the 2nd ACM International Workshop on Mobile Entity Localization and Tracking in GPS-less Environments*, pages 197—212, 2009. ACM, New York, NY, USA.

Cited on page: 67.

[15] A Barry, N. L Tye, and M. L Chang. Interactionless calendar-based training for 802.11 localization. *MASS '10: Proceedings of the 7th IEEE International Conference on Mobile Ad-hoc and Sensor Systems*, pages 166—194, 2010. IEEE, Washington, DC, USA.

Cited on page: 68.

[16] H Beadle, G Maguire, and M Smith. Using location and environment awareness in mobile communications. *ICICS '97: Proceedings of the 1st IEEE International Conference on Information, Communications and Signal Processing*, pages 1781—1785, 1997. IEEE, Washington, DC, USA.

Cited on page: 58.

[17] C Beckmann, S Consolvo, and A LaMarca. Some assembly required: Supporting end-user sensor installation in domestic ubiquitous computing environments. *UbiComp '04: Proceedings of the 6th International Conference on Ubiquitous Computing. Lecture Notes in Computer Science, 2004, Volume 3205*, pages 107—124, 2004. Springer, Heidelberg, Germany.

Cited on page: 51.

[18] U Blanke and B Schiele. Sensing location in the pocket. *UbiComp '08: Adjunct Proceedings of the 10th ACM International Conference on Ubiquitous Computing*, 2008. ACM, New York, NY, USA.

Cited on page: 63.

[19] P Bolliger. Redpin—adaptive, zero-configuration indoor localization through user collaboration. *MELT '08: Proceedings of the 1st ACM International Workshop on Mobile Entity Localization and Tracking in GPS-less Environments*, pages 55—60, 2008. ACM, New York, NY, USA.

Cited on page: 67.

[20] L Bonanni, E Arroyo, C.-H Lee, and T Selker. Smart sinks: real-world opportunities for context-aware interaction. *CHI '05: Proceedings of the 23rd ACM SIGCHI International Conference on Human Factors in Computing Systems (extended abstracts)*, 2005. ACM, New York, NY, USA.

Cited on page: 49.

[21] L Bonanni, M Hockenberry, D Zwarg, C Csikszentmihalyi, and H Ishii. Small business applications of Sourcemap: a web tool for sustainable design and supply chain transparency. *CHI '10: Proceedings of the 28th ACM SIGCHI International Conference on Human Factors in Computing Systems*, pages 937—946, 2010. ACM, New York, NY, USA.

Cited on page: 46.

[22] G Borriello, A Liu, T Offer, C Palistrant, and R Sharp. WALRUS: wireless acoustic location with room-level resolution using ultrasound. *MobiSys '05: Proceedings of the 3rd ACM International Conference on Mobile Systems, Applications, and Services*, pages 191—203, 2005. ACM, New York, NY, USA.

Cited on page: 64.

[23] R Bruno and F Delmastro. Design and analysis of a Bluetooth-based indoor localization system. *Personal Wireless Communications. Lecture Notes in Computer Science, 2003, Volume 2775*, pages 711—725. Springer, Heidelberg, Germany.

Cited on page: 72.

[24] J Burke, D Estrin, M Hansen, A Parker, N Ramanathan, S Reddy, and M. B Srivastava. Participatory sensing. *Proceedings of the Workshop on World-Sensor-Web, in conjunction with ACM SenSys '06*, pages 117—134.

Cited on page: 170.

[25] T Campbell, E Larson, G Cohn, R Alcaide, and S. N Patel. WATTR: a method for self-powered wireless sensing of water activity in the home. *UbiComp '10: Proceedings of the 12th ACM International Conference on Ubiquitous Computing*, pages 169—172, 2010. ACM, New York, NY, USA.

Cited on page: 49.

[26] P Castro, P Chiu, T Kremenek, and R. R Muntz. A probabilistic room location service for wireless networked environments. *UbiComp '01: Proceedings of the 3rd International Conference on Ubiquitous Computing. Lecture Notes in Computer Science, 2001, Volume 2201*, pages 18—34, 2001. Springer, Heidelberg, Germany.

Cited on page: 66.

[27] K Cheung, S. S Intille, and K Larson. An inexpensive Bluetooth-based indoor positioning hack. *UbiComp '06: Adjunct Proceedings of the 8th International Conference on Ubiquitous Computing*, 2006.

Cited on pages: 73 and 124.

[28] T Choudhury, S Consolvo, B. L Harrison, J Hightower, A LaMarca, L LeGrand, A Rahimi, A Rea, G Borriello, B Hemingway, P Klasnja, K Koscher, J. A Landay, J Lester, D Wyatt, and D Haehnel. The mobile sensing platform: An embedded activity recognition system. *IEEE Pervasive Computing*, 7(2):32—41, 2008.

Cited on page: 36.

[29] E. G Clary, M Snyder, R. D Ridge, J Copeland, A. A Stukas, J Haugen, and P Miene. Understanding and assessing the motivations of volunteers: a functional approach. *J Pers Soc Psychol*, 74(6):1516–30, May 1998.

Cited on page: 54.

[30] G Cohn, S Gupta, J Froehlich, E Larson, and S Patel. GasSense: Appliance-level, single-point sensing of gas activity in the home. *Pervasive '10: Proceedings of the 8th International Conference on Pervasive Computing. Lecture Notes in Computer Science, 2010, Volume 6030*, pages 265—282, 2010. Springer, Heidelberg, Germany.

Cited on page: 50.

[31] S Consolvo, K Everitt, I Smith, and J Landay. Design requirements for technologies that encourage physical activity. *CHI '06: Proceedings of the 24th ACM SIGCHI International Conference on Human factors in Computing Systems*, pages 457—466, 2006. ACM, New York, NY, USA.

Cited on page: 36.

[32] S Consolvo, P Klasnja, D McDonald, D Avrahami, J Froehlich, L LeGrand, R Libby, K Mosher, and J Landay. Flowers or a robot army?: encouraging awareness & activity with personal, mobile displays. *UbiComp '08: Proceedings of the 10th ACM*

International Conference on Ubiquitous Computing, pages 54—63, 2008. ACM, New York, NY, USA.

Cited on page: 36.

[33] S Consolvo, D McDonald, T Toscos, M Chen, J Froehlich, B. L Harrison, P Klasnja, A LaMarca, L LeGrand, R Libby, I Smith, and J Landay. Activity sensing in the wild: a field trial of UbiFit garden. *CHI '08: Proceeding of the 26th ACM SIGCHI International Conference on Human Factors in Computing Systems*, pages 1797—1806, 2008. ACM, New York, NY, USA.

Cited on page: 36.

[34] D Crawley, L Lawrie, F Winkelmann, W Buhl, Y Huang, C Pedersen, R Strand, R Liesen, D Fisher, M Witte, and J Glazer. Energyplus: creating a new-generation building energy simulation program. *Energy and Buildings*, 33(4):319—331, Dec 2001.

Cited on page: 95.

[35] D Crawley, J Hand, M Kummert, and B Griffith. Contrasting the capabilities of building energy performance simulation programs. *Building and Environment*, 43(4):661—673, 2008.

Cited on page: 95.

[36] A Dada, F. G von Reischach, and T Staake. Displaying dynamic carbon footprints of products on mobile phones. *Pervasive '08: Adjunct Proceedings of the 6th International Conference on Pervasive Computing*, 2008.

Cited on page: 46.

[37] S Darby. The effectiveness of feedback on energy consumption. *A Review for DEFRA of the Literature on Metering*, Dec 2006.

Cited on page: 32.

[38] J. J Davies, A. R Beresford, and A Hopper. Scalable, distributed, real-time map generation. *IEEE Pervasive Computing*, 5(4):47—54, Oct 2006.

Cited on page: 170.

[39] J. J Davies, D Cottingham, and B. D Jones. A sensor platform for sentient transportation research. *EuroSSC '06: Proceedings of the 1st European Conference on Smart Sensing and Context. Lecture Notes in Computer Science, 2006, Volume 4272*, pages 226—229, 2006. Springer, Heidelberg, Germany.

Cited on page: 47.

[40] S Dawson-Haggerty, X Jiang, G Tolle, J Ortiz, and D Culler. sMAP — a simple measurement and actuation profile for physical information. *SenSys '10: Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, pages 197—210, 2010. ACM, New York, NY, USA.

Cited on page: 75.

[41] S Dawson-Haggerty, J Ortiz, X Jiang, J Hsu, S Shankar, and D Culler. Enabling green building applications. *HotEmNets '10: Proceedings of the 6th Workshop on Hot Topics in Embedded Networked Sensors*, 2010.

Cited on page: 75.

[42] D. T Delaney, G. M. P O'Hare, and A Ruzzelli. Evaluation of energy-efficiency in lighting systems using sensor networks. *BuildSys '09: Proceedings of the 1st ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, in conjunction with SenSys '09*, pages 61—66, 2009. ACM, New York, NY, USA.

Cited on page: 93.

[43] A. K Dey and G. D Abowd. Towards a better understanding of context and context-awareness. *Workshop on The What, Who, Where, When, and How of Context-Awareness, in conjunction with CHI '00.*

Cited on pages: 116 and 117.

[44] T Dillahunt, G Becker, and J Mankoff. Motivating environmentally sustainable behavior changes with a virtual polar bear. *Pervasive '08 Workshop on Pervasive Persuasive Technology and Environmental Sustainability.*

Cited on page: 37.

[45] 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.

Cited on page: 55.

[46] P Dourish. What we talk about when we talk about context. *Pers Ubiquit Comput*, 8(1):19—30, Jan 2004.

Cited on page: 116.

[47] S Drenker and A Kader. Nonintrusive monitoring of electric loads. *IEEE Computer Applications in Power*, 12(4):47—51, 1999.

Cited on page: 43.

[48] P Dutta, M Feldmeier, J Paradiso, and D Culler. Energy metering for free: Augmenting switching regulators for real-time monitoring. *IPSN '08: Proceedings of the 7th IEEE International Conference on Information Processing in Sensor Networks*, pages 283—294, 2008. IEEE, Washington, DC, USA.

Cited on page: 43.

[49] M Eyole-Monono, R Harle, and A Hopper. POISE: An inexpensive, low-power location sensor based on electrostatics. *MobiQuitous '06: Proceedings of the 3rd IEEE International Conference on Mobile and Ubiquitous Systems: Networks and Services*, 2006. IEEE, Washington, DC, USA.

Cited on page: 56.

[50] S.-H Fang, J.-C Chen, H.-R Huang, and T.-N Lin. Metropolitan-scale location estimation using FM radio with analysis of measurements. *IWCMC '08: Proceedings*

of the International Wireless Communications and Mobile Computing Conference, pages 171—176, 2008.

Cited on page: 69.

[51] K. I Farkas, J Flinn, G Back, D Grunwald, and J Anderson. Quantifying the energy consumption of a pocket computer and a Java virtual machine. *SIGMETRICS '00: Proceedings of the 2000 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pages 252—263, 2000. ACM, New York, NY, USA.

Cited on page: 43.

[52] R Fielding and R Taylor. Principled design of the modern web architecture. *ACM Transactions on Internet Technology (TOIT)*, 2(2):115—150, May 2002.

Cited on pages: 74 and 156.

[53] C Fischer. Feedback on household electricity consumption: a tool for saving energy? *Energy Efficiency*, 1(1):79—104, 2008.

Cited on page: 32.

[54] G Fitzpatrick and G Smith. Technology-enabled feedback on domestic energy consumption: Articulating a set of design concerns. *IEEE Pervasive Computing*, 8(1):37—44.

Cited on page: 33.

[55] J Flinn and M Satyanarayanan. PowerScope: a tool for profiling the energy usage of mobile applications. *WMCSA '99: Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, pages 2—10, 1999. IEEE, Washington, DC, USA.

Cited on page: 43.

[56] J Fogarty, C Au, and S. E Hudson. Sensing from the basement: A feasibility study of unobtrusive and low-cost home activity recognition. *UIST 06: Proceedings of the ACM Symposium publisher = ACM, address = New York, NY, USA, year = 2006, on User Interface Software and Technology*, pages 91—100.

Cited on page: 49.

[57] B. J Fogg. Persuasive computers: perspectives and research directions. *CHI '98: Proceedings of the ACM SIGCHI International Conference on Human Factors in Computing Systems*, pages 225–232, 1998. ACM, New York, NY, USA.

Cited on page: 33.

[58] B. J Fogg. Persuasive technology: Using computers to change what we think and do. *Morgan Kaufman Publishers, San Francisco, CA, USA*, 2003.

Cited on pages: 33 and 37.

[59] R Fonseca, P Dutta, P Levis, and I Stoica. Quanto: Tracking energy in networked embedded systems. *OSDI '08: Proceedings of the 8th USENIX Symposium on*

Operating System Designs and Implementations, pages 323—338, 2008. USENIX, Berkeley, CA, USA.

Cited on page: 43.

[60] E Foxlin, M Harrington, and G Pfeifer. Constellation: a wide-range wireless motion-tracking system for augmented reality and virtual set applications. *SIGGRAPH '98: Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, pages 321—378, 1998.

Cited on page: 61.

[61] J Froehlich, T Dillahunt, P Klasnja, J Mankoff, S Consolvo, B. L Harrison, and J Landay. UbiGreen: investigating a mobile tool for tracking and supporting green transportation habits. *CHI '09: Proceedings of the 27th ACM SIGCHI International Conference on Human Factors in Computing Systems*, pages 1043—1052, 2009. ACM, New York, NY, USA.

Cited on page: 47.

[62] J Froehlich, K Everitt, J Fogarty, S. N Patel, and J Landay. Sensing opportunities for personalized feedback technology to reduce consumption. *CHI '09 Workshop: Defining the Role of HCI in the Challenges of Sustainability.*

Cited on page: 33.

[63] J Froehlich, L Findlater, and J Landay. The design of eco-feedback technology. *CHI '10: Proceedings of the 28th ACM SIGCHI International Conference on Human Factors in Computing Systems*, pages 1999—2008, 2010. ACM, New York, NY, USA.

Cited on page: 32.

[64] J Froehlich, E Larson, T Campbell, C Haggerty, J Fogarty, and S. N Patel. Hydrosense: infrastructure-mediated single-point sensing of whole-home water activity. *UbiComp '09: Proceedings of the 11th ACM International Conference on Ubiquitous Computing*, pages 235—244, 2009. ACM, New York, NY, USA.

Cited on page: 49.

[65] V Garg and N. K Bansal. Smart occupancy sensors to reduce energy consumption. *Energy and Buildings*, 32(1):81—87, Jun 2000.

Cited on page: 55.

[66] A Genco. Three step Bluetooth positioning. *LoCA '05: Proceedings of the International Workshop on Location and Context Awareness. Lecture Notes in Computer Science, 2005, Volume 3479*, pages 52—62, 2005. Springer, Heidelberg, Germany.

Cited on page: 74.

[67] A Giordano, M Chan, and H Habal. A novel location-based service and architecture. *PIMRC '95: Proceedings of the 6th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, pages 853—857, 1995. IEEE, Washington, DC, USA.

Cited on page: 69.

[68] U Gneezy, E Haruvy, and H Yafe. The inefficiency of splitting the bill. *The Economic Journal*, 114(495):265–280, 2004.

Cited on page: 80.

[69] R Gold. Optimal binary sequences for spread spectrum multiplexing (corresp.). *IEEE Transactions on Information Theory*, 13(4):619—621, Jan 1967.

Cited on page: 109.

[70] D Graumann, W Lara, J Hightower, and G Borriello. Real-world implementation of the location stack: the universal location framework. *WMCSA '04: Proceedings of the 5th IEEE Workshop on Mobile Computing Systems and Applications*, pages 122—128, 2004. IEEE, Washington, DC, USA.

Cited on page: 68.

[71] Y Gu and A Lo. A survey of indoor positioning systems for wireless personal networks. *IEEE Communications Surveys & Tutorials*, 11(1):13—32, 2009.

Cited on page: 57.

[72] D Guinard and V Trifa. Towards the web of things: Web mashups for embedded devices. *MEM '09: 2nd Workshop on Mashups, Enterprise Mashups and Lightweight Composition on the Web*, 2009.

Cited on pages: 41 and 74.

[73] D Guinard, M Weiss, and V Trifa. Are you energy-efficient? Sense it on the web! *Pervasive '09: Adjunct Proceedings of the 7th International Conference on Pervasive Computing*, 2009.

Cited on page: 41.

[74] S Gupta, M. S Reynolds, and S. N Patel. ElectriSense: single-point sensing using EMI for electrical event detection and classification in the home. *UbiComp '10: Proceedings of the 12th ACM International Conference on Ubiquitous Computing*, pages 139—148, 2010. ACM, New York, NY, USA.

Cited on page: 45.

[75] A Gustafsson and M Gyllenswärd. The power-aware cord: energy awareness through ambient information display. *CHI '05: Proceedings of the 23rd ACM SIGCHI International Conference on Human Factors in Computing Systems (extended abstracts)*, pages 1423—1426, 2005. ACM, New York, NY, USA.

Cited on page: 37.

[76] Y Gwon, R Jain, and T Kawahara. Robust indoor location estimation of stationary and mobile users. *INFOCOM '04: Proceedings of the 23rd Annual Joint Conference of the IEEE Computer and Communications Societies*, pages 1032—1043, 2004. IEEE, Washington, DC, USA.

Cited on pages: 73 and 134.

[77] A Haeberlen, E Flannery, A Ladd, A Rudys, D Wallach, and L Kavraki. Practical robust localization over large-scale 802.11 wireless networks. *MobiCom '04: Proceedings of the 10th Annual ACM/IEEE International Conference on Mobile Computing and Networking*, pages 70—84, 2004.

Cited on page: 66.

[78] M Haklay and P Weber. OpenStreetMap: User-generated street maps. *IEEE Pervasive Computing*, 7(4):12—18, Oct 2008.

Cited on page: 52.

[79] J Hallberg, M Nilsson, and K Synnes. Positioning with Bluetooth. *ICT '03: Proceedings of the 10th IEEE International Conference on Telecommunications*, pages 954—958, 2003. IEEE, Washington, DC, USA.

Cited on pages: 73 and 124.

[80] K.-H Han, S.-W Choi, B.-C Park, and J.-J Lee. An implementation of a wireless sensor network-based meter reading system. *SenSys '09: Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, pages 321—322, 2009. ACM, New York, NY, USA.

Cited on page: 50.

[81] G Hardin. The tragedy of the commons. *Science*, 162(5364):1243–8, Dec 1968.

Cited on page: 80.

[82] G. W Hart. Residential energy monitoring and computerized surveillance via utility power flows. *Technology and Society Magazine, IEEE*, 8(2):12—16, 1989.

Cited on pages: 43 and 156.

[83] G. W Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870—1891, 1992.

Cited on page: 43.

[84] A Harter and A Hopper. A distributed location system for the active office. *IEEE Network*, 8(1):62—70, 1994.

Cited on page: 58.

[85] A Harter, A Hopper, P Steggles, A Ward, and P Webster. The anatomy of a context-aware application. *MobiCom '99: Proceedings of the 5th Annual ACM International Conference on Mobile Computing and Networking*, pages 59—68, 1999. ACM, New York, NY, USA.

Cited on page: 59.

[86] S Hay. A global personal energy meter. *Pervasive '09: Adjunct Proceedings of the 7th International Conference on Pervasive Computing*, 2009.

Cited on page: 77.

[87] S Hay and R Harle. Bluetooth tracking without discoverability. *LoCA '09: Proceedings of the 4th International Symposium on Location and Context Awareness. Lecture Notes in Computer Science, 2009, Volume 5561*, pages 120—137, 2009. Springer, Heidelberg, Germany.

Cited on pages: 18, 19, 115, 122, 131, 135, and 137.

[88] S Hay, J. F Newman, and A Rice. Sentient computing meets social networking. *W3C Workshop on the Future of Social Networking 2009.*

Cited on page: 33.

[89] S Hay, S. T Rassia, and A. R Beresford. Estimating personal energy expenditure with location data. *PerHealth '10: Proceedings of the 1st IEEE PerCom Workshop on Pervasive Healthcare*, 2010. IEEE, Washington, DC, USA.

Cited on page: 139.

[90] S Hay and A Rice. The case for apportionment. *BuildSys '09: Proceedings of the 1st ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, in conjunction with SenSys '09*, pages 13—18, 2009. ACM, New York, NY, USA.

Cited on page: 77.

[91] M Hazas and A Ward. A novel broadband ultrasonic location system. *UbiComp '02: Proceedings of the 4th International Conference on Ubiquitous Computing. Lecture Notes in Computer Science, 2002, Volume 2498*, pages 264—280, 2002. Springer, Heidelberg, Germany.

Cited on page: 60.

[92] M Hazas and A Ward. A high performance privacy-oriented location system. *PerCom '03: Proceedings of the 1st Annual IEEE International Conference on Pervasive Computing and Communications*, pages 216—223, 2003. IEEE, Washington, DC, USA.

Cited on page: 60.

[93] J Hightower and G Borriello. Location sensing techniques. Aug 2001.

Cited on page: 57.

[94] J Hightower and G Borriello. Location systems for ubiquitous computing. *Computer*, 34(8):57—66, Dec 2001.

Cited on page: 57.

[95] T. G Holmes. Eco-visualization: combining art and technology to reduce energy consumption. *Creativity and Cognition '07: Proceedings of the 6th ACM SIGGCHI Conference on Creativity & Cognition*, pages 153—162, 2007. ACM, New York, NY, USA.

Cited on page: 38.

[96] A Hopper. Sentient computing? *Computer Systems: Theory, Technology, and Applications: A Tribute to Roger Needham, Series Monographs in Computer Science*, pages 125—131, 2003.

Cited on page: 26.

[97] A Hopper. Opinion first person—if you ask me—computing the planet's future. *Engineering & Technology*, 2(7):24, Jul 2007.

Cited on pages: 24 and 145.

[98] 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.

Cited on pages: 25 and 34.

[99] J Hsu, P Mohan, X Jiang, J Ortiz, S Shankar, S Dawson-Haggerty, and D Culler. HBCI: Human-building-computer interaction. *BuildSys '10: Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, in conjunction with SenSys '10*, pages 55—60, 2010. ACM, New York, NY, USA.

Cited on page: 75.

[100] A Huang. The use of Bluetooth in Linux and location aware computing. *Master of Science dissertation, Massachusetts Institute of Technology*, 2005.

Cited on pages: 72 and 124.

[101] A Hylick, R Sohan, A Rice, and B. D Jones. An analysis of hard drive energy consumption. *MASCOTS '08: Proceedings of the 16th Annual IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*, pages 103—112, 2008. IEEE, Washington, DC, USA.

Cited on page: 104.

[102] F Ichikawa, J Chipchase, and R Grignani. Where's the phone? a study of mobile phone location in public spaces. *Mobility '05: Proceedings of the 2nd IEEE International Conference on Mobile Technology, Applications and Systems, 2005*, pages 797—804, 2005. IEEE, Washington, DC, USA.

Cited on page: 63.

[103] M Jevring, R de Groote, and C Hesselman. Dynamic optimization of Bluetooth networks for indoor localization. *AASN '08: 1st International Workshop on Automated and Autonomous Sensor Networks*, 2008.

Cited on page: 145.

[104] X Jiang, S Dawson-Haggerty, P Dutta, and D Culler. Design and implementation of a high-fidelity AC metering network. *IPSN '09: Proceedings of the 8th ACM International Conference on Information Processing in Sensor Networks*, pages 253—264, 2009. ACM, New York, NY, USA.

Cited on pages: 42 and 108.

[105] X Jiang, M. V Ly, J Taneja, P Dutta, and D Culler. Experiences with a high-fidelity wireless building energy auditing network. *SenSys '09: Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, pages 113—126, 2009. ACM, New York, NY, USA.

Cited on page: 44.

[106] C Jin and K Fujinami. Is my mobile phone in the chest pocket?: Knowing mobile phone's location on the body. *Pervasive '09: Adjunct Proceedings of the 7th International Conference on Pervasive Computing*, pages 189—192, 2009.

Cited on page: 63.

[107] D Jung and A Savvides. Estimating building consumption breakdowns using on/off state sensing and incremental sub-meter deployment. *SenSys '10: Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, pages 225—238, 2010. ACM, New York, NY, USA.

Cited on page: 45.

[108] K Kaemarungsi. Distribution of WLAN received signal strength indication for indoor location determination. *ISWPC '06: Proceedings of the 1st IEEE International Symposium on Wireless Pervasive Computing*, pages 1—6, 2006. IEEE, Washington, DC, USA.

Cited on page: 69.

[109] K Kaemarungsi and P Krishnamurthy. Properties of indoor received signal strength for WLAN location fingerprinting. *MobiQuitous '04: Proceedings of the 1st IEEE International Conference on Mobile and Ubiquitous Systems: Networks and Services*, pages 14—23, 2004. IEEE, Washington, DC, USA.

Cited on page: 69.

[110] K Kappel and T Grechenig. "show-me": water consumption at a glance to promote water conservation in the shower. *Persuasive '09: Proceedings of the 4th International Conference on Persuasive Technology*, page 26, 2009.

Cited on page: 49.

[111] F Kawsar, T Nakajima, and K Fujinami. Deploy spontaneously: supporting end-users in building and enhancing a smart home. *UbiComp '08: Proceedings of the 10th ACM International Conference on Ubiquitous Computing*, pages 282—291, 2008. ACM, New York, NY, USA.

Cited on page: 52.

[112] M Kazandjieva, O Gnawali, B Heller, P Levis, and C Kozyrakis. Identifying energy waste through dense power sensing and utilization monitoring. *Technical Report SING-10-03, Stanford University*, 2010.

Cited on page: 42.

[113] M. A Kazandjieva, B Heller, P Levis, and C Kozyrakis. Energy dumpster diving. *HotPower '09: Proceedings of the 2nd Workshop on Power Aware Computing*, 2009.

Cited on page: 42.

[114] J.-W Kim, Y.-K Kim, and T.-J Nam. The ténéré: design for supporting energy conservation behaviors. *CHI '09: Proceedings of the 27th ACM SIGCHI International Conference on Human Factors in Computing Systems (extended abstracts)*, pages 2643—2646, 2009. ACM, New York, NY, USA.

Cited on page: 38.

[115] Y Kim, T Schmid, Z Charbiwala, J Friedman, and M. B Srivastava. NAWMS: nonintrusive autonomous water monitoring system. *SenSys '08: Proceedings of the 6th ACM Conference on Embedded Networked Sensor Systems*, pages 309—322, 2008. ACM, New York, NY, USA.

Cited on page: 48.

[116] Y Kim, T Schmid, Z Charbiwala, and M. B Srivastava. Viridiscope: design and implementation of a fine grained power monitoring system for homes. *UbiComp '09: Proceedings of the 11th ACM International Conference on Ubiquitous Computing*, pages 245—254, 2009. ACM, New York, NY, USA.

Cited on page: 45.

[117] Y Kim, T Schmid, M. B Srivastava, and Y Wang. Challenges in resource monitoring for residential spaces. *BuildSys '09: Proceedings of the 1st ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, in conjunction with SenSys '09*, pages 1—7, 2009. ACM, New York, NY, USA.

Cited on pages: 48 and 49.

[118] D Kirsch and T Starner. The locust swarm: An environmentally-powered, networkless location and messaging system. *ISWC '97: Proceedings of the 1st IEEE International Symposium on Wearable Computers*, pages 169—170, 1997. IEEE, Washington, DC, USA.

Cited on pages: 59 and 93.

[119] J Kitzes, M Wackernagel, J Loh, A Peller, S Goldfinger, D Cheng, and K Tea. Shrink and share: humanity's present and future ecological footprint. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1491):467—475, Jul 2007.

Cited on page: 23.

[120] J Kleissl and Y Agarwal. Cyber-physical energy systems: Focus on smart buildings. *DAC '10: Proceedings of the 47th ACM Design Automation Conference*, pages 749—754, 2010. ACM, New York, NY, USA.

Cited on page: 39.

[121] R Kohlenberg, T Phillips, and W Proctor. A behavioral analysis of peaking in residential electrical-energy consumers. *J Appl Behav Anal*, 9(1):13–18, Mar 1976.

Cited on page: 32.

[122] A Kotanen, M Hännikäinen, H Leppakoski, and T. D Hamalainen. Experiments on local positioning with Bluetooth. *ITCC '03: Proceedings of the IEEE International Conference on Information Technology: Computers and Communications*, pages 297—303, 2003. IEEE, Washington, DC, USA.

Cited on page: 73.

[123] M Kranz, C Fischer, and A Schmidt. A comparative study of DECT and WLAN signals for indoor localization. *PerCom '10: Proceedings of the 8th Annual IEEE International Conference on Pervasive Computing and Communications*, pages 235—243, 2010. IEEE, Washington, DC, USA.

Cited on page: 71.

[124] J Krumm, G Cermak, and E Horvitz. RightSPOT: A novel sense of location for a smart personal object. *UbiComp '03: Proceedings of the 5th International Conference on Ubiquitous Computing. Lecture Notes in Computer Science, 2003, Volume 2864*, pages 36—43, 2003. Springer, Heidelberg, Germany.

Cited on page: 69.

[125] J Krumm, S Harris, B Meyers, B Brumitt, M Hale, and S Shafer. Multi-camera multi-person tracking for EasyLiving. *VS '00: Proceedings of the 3rd IEEE International Workshop on Visual Surveillance*, pages 3—10, 2000. IEEE, Washington, DC, USA.

Cited on page: 57.

[126] J Krumm, L Williams, and G Smith. SmartMoveX on a graph - an inexpensive active badge tracker. *UbiComp '02: Proceedings of the 4th International Conference on Ubiquitous Computing. Lecture Notes in Computer Science, 2002, Volume 2498*, pages 299—307, 2002. Springer, Heidelberg, Germany.

Cited on page: 61.

[127] K Kunze, P Lukowicz, H Junker, and G Tröster. Where am I: recognizing on-body positions of wearable sensors. *LoCA '05: Proceedings of the International Workshop on Location and Context Awareness. Lecture Notes in Computer Science, 2005, Volume 3479*, pages 264—275, 2005. Springer, Heidelberg, Germany.

Cited on page: 63.

[128] P Kurp. Green computing. *Communications of the ACM*, 51(10):11—13, Oct 2008.

Cited on page: 24.

[129] A LaMarca, Y Chawathe, S Consolvo, J Hightower, I Smith, J Scott, T Sohn, J Howard, J Hughes, F Potter, J Tabert, P Powledge, G Borriello, and B. N Schilit. Place Lab: Device positioning using radio beacons in the wild. *Pervasive '05: Proceedings of the 3rd International Conference on Pervasive Computing. Lecture Notes in Computer Science, 2005, Volume 3468*, pages 116—133, 2005. Springer, Heidelberg, Germany.

Cited on pages: 68 and 117.

[130] C Laughman, K Lee, R Cox, S Shaw, S Leeb, L Norford, and P Armstrong. Power signature analysis. *IEEE Power and Energy Magazine*, 1(2):56—63, 2003.

Cited on page: 44.

[131] D Lee, T Kato, H Cho, T Toyomura, T Yamazaki, and M Hahn. Bit-watt system: Information-energy integrated system towards assistive home services. *Pervasive '09: Adjunct Proceedings of the 7th International Conference on Pervasive Computing*, 2009.

Cited on page: 41.

[132] J Lifton, M Feldmeier, Y Ono, C Lewis, and J Paradiso. A platform for ubiquitous sensor deployment in occupational and domestic environments. *IPSN '07: Proceedings of the 6th ACM International Conference on Information Processing in Sensor Networks*, pages 119—127, 2007. ACM, New York, NY, USA.

Cited on page: 41.

[133] J Lin, L Mamykina, S Lindtner, G Delajoux, and H. B Strub. Fish'n'steps: encouraging physical activity with an interactive computer game. *UbiComp '06: Proceedings of the 8th International Conference on Ubiquitous Computing. Lecture Notes in Computer Science, 2006, Volume 4206*, pages 261—278, 2006. Springer, Heidelberg, Germany.

Cited on page: 35.

[134] T Lovett, E O'Neill, J Irwin, and D Pollington. The calendar as a sensor: analysis and improvement using data fusion with social networks and location. *UbiComp '10: Proceedings of the 12th ACM International Conference on Ubiquitous Computing*, pages 3—12, 2010. ACM, New York, NY, USA.

Cited on page: 117.

[135] L Lutzenhiser. Social and behavioral aspects of energy use. *Annual Review of Energy and the Environment*, 18:247—289, 1993.

Cited on page: 80.

[136] D. J. C MacKay. Sustainable energy - without the hot air. *UIT, Cambridge, UK.*

Cited on pages: 47, 78, 98, and 161.

[137] A Madhavapeddy and A Tse. A study of Bluetooth propagation using accurate indoor location mapping. *UbiComp '05: Proceedings of the 7th International Conference on Ubiquitous Computing. Lecture Notes in Computer Science, 2005, Volume 3660*, pages 105—122, 2005. Springer, Heidelberg, Germany.

Cited on pages: 72, 74, and 135.

[138] A Mandal, C Lopes, T Givargis, A Haghighat, R Jurdak, and P Baldi. Beep: 3D indoor positioning using audible sound. *CCNC '05: Proceedings of the 2nd IEEE Consumer Communications and Networking Conference*, pages 348—353, 2005. IEEE, Washington, DC, USA.

Cited on page: 63.

[139] J Mankoff, S. R Fussell, T Dillahunt, R Glaves, C Grevet, M Johnson, D Matthews, H. S Matthews, R McGuire, and R Thompson. Stepgreen.org: Increasing energy saving behaviors via social networks. *ICWSM '10: Proceedings of the 4th International AAAI Conference on Weblogs and Social Media*, 2010.

Cited on page: 37.

[140] J Mankoff, D Matthews, S Fussell, and M Johnson. Leveraging social networks to motivate individuals to reduce their ecological footprints. *HICSS '07: Proceedings of the 40th IEEE Hawaii International Conference on System Sciences*, page 87, 2007. IEEE, Washington, DC, USA.

Cited on page: 37.

[141] K Mansley, A. R Beresford, and D Scott. The carrot approach: Encouraging use of location systems. *UbiComp '04: Proceedings of the 6th International Conference on Ubiquitous Computing. Lecture Notes in Computer Science, 2004, Volume 3205*, pages 366—383, 2004. Springer, Heidelberg, Germany.

Cited on pages: 119 and 170.

[142] A Matic, A Papliatseyeu, V Osmani, and O Mayora-Ibarra. Tuning to your position: FM radio based indoor localization with spontaneous recalibration. *PerCom '10: Proceedings of the 8th Annual IEEE International Conference on Pervasive Computing and Communications*, pages 153—161, 2010. IEEE, Washington, DC, USA.

Cited on pages: 66, 67, and 70.

[143] J McCarthy and E. S Meidel. ACTIVEMAP: A visualization tool for location awareness to support informal interactions. *HUC '99: Proceedings of the International Symposium on Handheld and Ubiquitous Computing. Lecture Notes in Computer Science, 1999, Volume 1707*, pages 158—170, 1999. Springer, Heidelberg, Germany.

Cited on page: 33.

[144] E Miluzzo, N. D Lane, K Fodor, R Peterson, H Lu, M Musolesi, S. B Eisenman, X Zheng, and A. T Campbell. Sensing meets mobile social networks: The design, implementation and evaluation of the CenceMe application. *SenSys '08: Proceedings of the 6th ACM Conference on Embedded Networked Sensor Systems*, pages 337—350, 2008. ACM, New York, NY, USA.

Cited on page: 34.

[145] P Mohan, V Padmanabhan, and R Ramjee. Nericell: rich monitoring of road and traffic conditions using mobile smartphones. *SenSys '08: Proceedings of the 6th ACM Conference on Embedded Networked Sensor Systems*, pages 323—336, 2008. ACM, New York, NY, USA.

Cited on page: 43.

[146] I Mulder, B Hulsebosch, G Lenzini, and M Bargh. Reading the tea-leaves in an intelligent coffee corner: understanding behavior by using sensory data. *Measuring Behaviour '08: Proceedings of the 6th International Conference on Methods and Techniques in Behavioural Research*, pages 68—71, 2008.

Cited on page: 33.

[147] M Mun, S Reddy, K Shilton, N Yau, J Burke, D Estrin, M Hansen, E Howard, R West, and P Boda. PEIR, the personal environmental impact report, as a platform for participatory sensing systems research. *MobiSys '09: Proceedings of the 7th ACM International Conference on Mobile Systems, Applications, and Services*, pages 55—68, 2009. ACM, New York, NY, USA.

Cited on pages: 33 and 170.

[148] J Nagle. Congestion Control in IP/TCP Internetworks. RFC 896, January 1984.

Cited on page: 177.

[149] F Naya, H Noma, R Ohmura, and K Kogure. Bluetooth-based indoor proximity sensing for nursing context awareness. *ISWC '05: Proceedings of the 9th IEEE International Symposium on Wearable Computers*, pages 212—213, 2005. IEEE, Washington, DC, USA.

Cited on page: 73.

[150] Y Nishida, H Aizawa, T Hori, N. H Hoffman, T Kanade, and M Kakikura. 3D ultrasonic tagging system for observing human activity. *IROS '03: Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 785—791, 2003.

Cited on page: 145.

[151] M Nottingham and R Sayre. The Atom Syndication Format. RFC 4287 (Proposed Standard), December 2005.

Cited on page: 156.

[152] O Nov. What motivates Wikipedians? *Communications of the ACM*, 50(11):60—64, Oct 2007.

Cited on pages: 52 and 54.

[153] M Ocana, L Bergasa, M Sotelo, J Nuevo, and R Flores. Indoor robot localization system using WiFi signal measure and minimizing calibration effort. *ISIE '05: Proceedings of the IEEE International Symposium on Industrial Electronics*, pages 1545—1550, 2005. IEEE, Washington, DC, USA.

Cited on page: 66.

[154] R. J Orr, G. D Abowd, C Atkeson, I Essa, and R Gregor. The smart floor: A mechanism for natural user identification and tracking. *CHI '00: Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, pages 275—276, 2000. ACM, New York, NY, USA.

Cited on page: 56.

[155] V Otsason, A Varshavsky, A LaMarca, and E de Lara. Accurate GSM indoor localization. *UbiComp '05: Proceedings of the 7th International Conference on Ubiquitous Computing. Lecture Notes in Computer Science, 2005, Volume 3660*, pages 141—158, 2005. Springer, Heidelberg, Germany.

Cited on page: 68.

[156] M. S Owen, editor. *ASHRAE Handbook–Fundamentals.* American Society of Heating, Refrigerating and Air-Conditioning Engineers, 2009.

Cited on page: 98.

[157] D Pandya, R Jain, and E Lupu. Indoor location estimation using multiple wireless technologies. *PIMRC '03: 14th IEEE Proceedings on Personal, Indoor and Mobile Radio Communications*, pages 2208—2212, 2003. IEEE, Washington, DC, USA.

Cited on page: 73.

[158] A Papliatseyeu, V Osmani, and O Mayora-Ibarra. FINDR low-cost indoor positioning using FM radio. *Mobile Wireless Middleware, Operating Systems and Applications. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, 2009, Volume 7*, pages 15—26.

Cited on page: 70.

[159] J Paradiso, C Abler, K yuh Hsiao, and M. S Reynolds. The magic carpet: physical sensing for immersive environments. *CHI '97: Proceedings of the ACM SIGCHI International Conference on Human Factors in Computing Systems (extended abstracts)*, pages 227—278, 1997. ACM, New York, NY, USA.

Cited on page: 56.

[160] J Paradiso and V Speakers. Some novel applications for wireless inertial sensors. *Nanotech '06: Proceedings of the 9th Annual NSTI Nanotechnology Conference*, pages 431—434, 2006.

Cited on page: 41.

[161] S. N Patel, S Gupta, and M. S Reynolds. The design and evaluation of an end-user-deployable, whole house, contactless power consumption sensor. *CHI '10: Proceedings of the 28th ACM SIGCHI International Conference on Human Factors in Computing Systems*, pages 2471—2480, 2010. ACM, New York, NY, USA.

Cited on page: 39.

[162] S. N Patel, J. A Kientz, G. R Hayes, S Bhat, and G. D Abowd. Farther than you may think: An empirical investigation of the proximity of users to their mobile phones. *UbiComp '06: Proceedings of the 8th International Conference on Ubiquitous Computing. Lecture Notes in Computer Science, 2006, Volume 4206*, pages 123—140, 2006. Springer, Heidelberg, Germany.

Cited on page: 119.

[163] S. N Patel, M. S Reynolds, and G. D Abowd. Detecting human movement by differential air pressure sensing in HVAC system ductwork: An exploration in infrastructure mediated sensing. *Pervasive '08: Proceedings of the 6th International Conference on Pervasive Computing. Lecture Notes in Computer Science, 2008, Volume 5013*, pages 1—18, 2008. Springer, Heidelberg, Germany.

Cited on page: 65.

[164] S. N Patel, T Robertson, J. A Kientz, M. S Reynolds, and G. D Abowd. At the flick of a switch: Detecting and classifying unique electrical events on the residential power line. *UbiComp '07: Proceedings of the 9th International Conference on Ubiquitous Computing. Lecture Notes in Computer Science, 2007, Volume 4717*, pages 271—288, 2007. Springer, Heidelberg, Germany.

Cited on page: 45.

[165] S. N Patel, K. N Truong, and G. D Abowd. Powerline positioning: A practical sub-room-level indoor location system for domestic use. *UbiComp '06: Proceedings of the 8th International Conference on Ubiquitous Computing. Lecture Notes in Computer Science, 2006, Volume 4206*, pages 441—458, 2006. Springer, Heidelberg, Germany.

Cited on page: 64.

[166] T Pering, Y Agarwal, R. K Gupta, and R Want. Coolspots: reducing the power consumption of wireless mobile devices with multiple radio interfaces. *MobiSys '06: Proceedings of the 4th ACM International Conference on Mobile Systems, Applications, and Services*, pages 220—232, 2006. ACM, New York, NY, USA.

Cited on page: 43.

[167] 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*, 8(1):16—33, 2007.

Cited on page: 32.

[168] D Plummer. Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware. RFC 826 (Standard), November 1982. Updated by RFCs 5227, 5494.

Cited on pages: 170 and 172.

[169] N. B Priyantha, A Chakraborty, and H Balakrishnan. The Cricket location-support system. *MobiCom '00: Proceedings of the 6th Annual ACM International Conference on Mobile Computing and Networking*, pages 32—43, 2000. ACM, New York, NY, USA.

Cited on page: 60.

[170] N. B Priyantha, A Miu, H Balakrishnan, and S Teller. The cricket compass for context-aware mobile applications. *MobiCom '01: Proceedings of the 7th Annual ACM International Conference on Mobile Computing and Networking*, (1–14), 2001. ACM, New York, NY, USA.

Cited on page: 60.

[171] C Randell and H Muller. Low cost indoor positioning system. *UbiComp '01: Proceedings of the 3rd International Conference on Ubiquitous Computing. Lecture Notes in Computer Science, 2001, Volume 2201*, pages 42—48, 2001. Springer, Heidelberg, Germany.

Cited on page: 61.

[172] S Reddy, J Burke, D Estrin, M Hansen, and M. B Srivastava. Determining transportation mode on mobile phones. *ISWC '08: Proceedings of the 12th IEEE International Symposium on Wearable Computers*, pages 25—28, 2008. IEEE, Washington, DC, USA.

Cited on page: 47.

[173] I Reichart and R Hischier. The environmental impact of getting the news. *Journal of Industrial Ecology*, 6(3-4):185—200.

Cited on page: 46.

[174] A Rice. Dependable systems for Sentient Computing. Technical Report UCAM-CL-TR-686, University of Cambridge, Computer Laboratory, May 2007.

Cited on page: 57.

[175] A Rice and S Hay. Decomposing power measurements for mobile devices. *PerCom '10: Proceedings of the 8th Annual IEEE International Conference on Pervasive Computing and Communications*, pages 70—78, 2010. IEEE, Washington, DC, USA.

Cited on pages: 91 and 169.

[176] A Rice and S Hay. Measuring mobile phone energy consumption for 802.11 wireless networking. *Pervasive and Mobile Computing*, 6(6):593–606, Dec 2010.

Cited on pages: 18, 19, 91, 109, 110, 111, 147, 169, 171, 172, 173, 174, 175, 176, 177, and 178.

[177] A Rice, S Hay, and D Ryder-Cook. A limited-data model of building energy consumption. *BuildSys '10: Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, in conjunction with SenSys '10*, pages 67—72, 2010. ACM, New York, NY, USA.

Cited on pages: 17, 18, 91, 97, 99, 100, 101, and 103.

[178] A Rice and O. J Woodman. Crowd-sourcing world models with OpenRoomMap. *PerCom '10: Adjunct Proceedings of the 8th Annual IEEE International Conference on Pervasive Computing and Communications*, 2010. IEEE, Washington, DC, USA.

Cited on pages: 17, 53, 54, and 93.

[179] 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.

Cited on page: 55.

[180] T Richardson, Q Stafford-Fraser, K Wood, and A Hopper. Virtual network computing. *IEEE Internet Computing*, 2(1):33—38, Jan 1998.

Cited on page: 104.

[181] Y Rogers, W. R Hazlewood, P Marshall, N Dalton, and S Hertrich. Ambient influence: Can twinkly lights lure and abstract representations trigger behavioral change? *UbiComp '10: Proceedings of the 12th ACM International Conference on Ubiquitous Computing*, pages 261—270, 2010. ACM, New York, NY, USA.

Cited on page: 35.

[182] A Rowe, M Berges, and R Rajkumar. Contactless sensing of appliance state transitions through variations in electromagnetic fields. *BuildSys '10: Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, in conjunction with SenSys '10*, pages 19—24, 2010. ACM, New York, NY, USA.

Cited on page: 45.

[183] B. N Schilit, D Hilbert, and J Trevor. Context-aware communication. *IEEE Wireless Communications*, 9(5):46—54, 2002.

Cited on page: 58.

[184] A Schwaighofer, M Grigoras, V Tresp, and C Hoffmann. GPPS: A Gaussian process positioning system. *Advances in Neural Information Processing Systems*, 16:579—-586, Feb 2004.

Cited on pages: 66 and 71.

[185] J Scott and M Hazas. User-friendly surveying techniques for location-aware systems. *UbiComp '03: Proceedings of the 5th International Conference on Ubiquitous Computing. Lecture Notes in Computer Science, 2003, Volume 2864*, pages 44—53, 2003. Springer, Heidelberg, Germany.

Cited on page: 59.

[186] A Shah, T Christian, C Patel, C Bash, and R Sharma. Assessing ICT's environmental impact. *Computer*, 42(7):91—93, Jul 2009.

Cited on page: 46.

[187] M Shiraishi, Y Washio, C Takayama, V Lehdonvirta, H Kimura, and T Nakajima. Using individual, social and economic persuasion techniques to reduce CO2 emissions in a family setting. *Persuasive '09: Proceedings of the 4th International Conference on Persuasive Technology*, page 13, 2009.

Cited on page: 37.

[188] Q Stafford-Fraser. On site: The life and times of the first web cam. *Communications of the ACM*, 44(7):25—26, 2001.

Cited on page: 116.

[189] U Steinhoff and B Schiele. Dead reckoning from the pocket—an experimental study. *PerCom '10: Proceedings of the 8th Annual IEEE International Conference on Pervasive Computing and Communications*, pages 162—170, 2010. IEEE, Washington, DC, USA.

Cited on page: 63.

[190] E Stuntebeck, S. N Patel, T Robertson, M. S Reynolds, and G. D Abowd. Wideband powerline positioning for indoor localization. *UbiComp '08: Proceedings of the 10th ACM International Conference on Ubiquitous Computing*, pages 94—103, 2008. ACM, New York, NY, USA.

Cited on page: 65.

[191] S Taherian, M Pias, G Coulouris, and J Crowcroft. Profiling energy use in households and office spaces. *e-Energy '10: Proceedings of the 1st ACM International Conference on Energy-Efficient Computing and Networking*, pages 21—30, 2010. ACM, New York, NY, USA.

Cited on page: 40.

[192] S Tarzia, R Dick, P Dinda, and G Memik. Sonar-based measurement of user presence and attention. *UbiComp '09: Proceedings of the 11th ACM International Conference on Ubiquitous Computing*, pages 89—92, 2009. ACM, New York, NY, USA.

Cited on page: 56.

[193] A Thiagarajan, J Biagioni, T Gerlich, and J Eriksson. Cooperative transit tracking using smart-phones. *SenSys '10: Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, pages 85—98, 2010. ACM, New York, NY, USA.

Cited on page: 47.

[194] T Toscos, A Faber, S An, and M Gandhi. Chick clique: persuasive technology to motivate teenage girls to exercise. *CHI '06: Proceedings of the 24th ACM SIGCHI International Conference on Human factors in Computing Systems (extended abstracts)*, pages 1873—1878, 2006. ACM, New York, NY, USA.

Cited on page: 35.

[195] A Varshavsky, E de Lara, J Hightower, A LaMarca, and V Otsason. GSM indoor localization. *Pervasive and Mobile Computing*, 3(6):698—720, 2007.

Cited on page: 68.

[196] D Wang, C Federspiel, and F Rubinstein. Modeling occupancy in single person offices. *Energy and Buildings*, 37(2):121—126, Feb 2005.

Cited on page: 55.

[197] R Want and A Hopper. Active badges and personal interactive computing objects. *IEEE Transactions on Consumer Electronics*, 38(1):10—20, Feb 1992.

Cited on page: 58.

[198] R Want, A Hopper, V Falcao, and J Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10(1):91–102, 1992.

Cited on page: 57.

[199] R Want, T Pering, G Borriello, and K. I Farkas. Disappearing hardware. *IEEE Pervasive Computing*, 1(1):36—47, 2002.

Cited on page: 41.

[200] R Want, B. N Schilit, N. I Adams, R Gold, K Petersen, D Goldberg, J. R Ellis, and M Weiser. An overview of the ParcTab ubiquitous computing experiment. *IEEE Personal Communications*, 2(6):28–33, Jun 1995.

Cited on page: 58.

[201] A Ward, A Jones, and A Hopper. A new location technique for the active office. *IEEE Personal Communications*, 4(5):42—47, Oct 1997.

Cited on page: 59.

[202] M Weiser. The computer for the 21st century. *ACM SIGMOBILE Mobile Computing and Communications Review*, 3(3):3—11, Jul 1999.

Cited on page: 104.

[203] M Weiss, T Graml, T Staake, F Mattern, and E Fleisch. Handy feedback: Connecting smart meters with mobile phones. *MUM '09: Proceedings of the ACM*

International Conference on Mobile and Ubiquitous Mulitimedia, page 15, 2009. ACM, New York, NY, USA.

Cited on page: 39.

[204] R. A Winett, J. H Kagel, R. C Battalio, and R. C Winkler. Effects of monetary rebates, feedback, and information on residential electricity conservation. *Journal of Applied Psychology*, 63(1):73—80, Dec 1978.

Cited on page: 32.

[205] G Wood and M Newborough. Energy-use information transfer for intelligent homes: Enabling energy conservation with central and local displays. *Energy and Buildings*, 39(4):495–503, Dec 2007.

Cited on page: 26.

[206] O. J Woodman. Pedestrian localisation for indoor environments. *Ph.D. dissertation, University of Cambridge*, 2010.

Cited on page: 62.

[207] O. J Woodman and R Harle. Pedestrian localisation for indoor environments. *UbiComp '08: Proceedings of the 10th ACM International Conference on Ubiquitous Computing*, pages 114—123, 2008. ACM, New York, NY, USA.

Cited on page: 62.

[208] O. J Woodman and R Harle. RF-based initialisation for inertial pedestrian tracking. *Pervasive '09: Proceedings of the 7th International Conference on Pervasive Computing. Lecture Notes in Computer Science, 2009, Volume 5538*, pages 238—255, 2009. Springer, Heidelberg, Germany.

Cited on page: 67.

[209] A Youssef, J Krumm, E Miller, G Cermak, and E Horvitz. Computing location from ambient FM radio signals. *WCNC '05: Proceedings of the 2005 IEEE Conference on Wireless Communications and Networking Conference*, pages 824—829, 2005. IEEE, Washington, DC, USA.

Cited on page: 69.

[210] M Youssef and A Agrawala. The Horus WLAN location determination system. *MobiSys '05: Proceedings of the 3rd ACM International Conference on Mobile Systems, Applications, and Services*, pages 205—218, 2005. ACM, New York, NY, USA.

Cited on page: 68.

[211] Y Zheng, Q Li, Y Chen, X Xie, and W.-Y Ma. Understanding mobility based on GPS data. *UbiComp '08: Proceedings of the 10th ACM International Conference on Ubiquitous Computing*, pages 312—321, 2008. ACM, New York, NY, USA.

Cited on page: 47.

[212] Y Zheng, L Liu, L Wang, and X Xie. Learning transportation mode from raw GPS data for geographic applications on the web. *WWW '08: Proceeding of the 17th*

ACM International Conference on the World Wide Web, pages 247—250, 2008. ACM, New York, NY, USA.

Cited on page: 47.