A Global Personal Energy Meter

Simon Hay

University of Cambridge Computer Laboratory

Abstract. 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. I envisage a personal energy meter which can record and apportion an individual's energy usage in order to provide baseline information and incentives for reducing the environmental impact of our lives. Contextual information will be crucial for apportioning the use and energy costs of shared resources. In order to obtain this it will be necessary to develop low cost, low infrastructure location systems that can be deployed on a truly global scale.

1 Introduction

Controlling our 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 [1].

I envisage a personal energy meter (PEM) which can record and apportion an individual's energy usage. This is just one of the roles which computing might play in ensuring the future of our planet [2]. It could provide a mechanism to raise awareness of our consumption and help us reduce our impact [3]. This would depend on a global sensor network and poses a number of challenges for pervasive computing.

The Swiss Société 2000 Watts project aims to reduce the average overall rate of consumption of each citizen to 2000 W without lowering their standard of living [4]. Unfortunately, it is at present very difficult for an individual to assess whether or not he is achieving that target: to judge the success of a weight loss programme one requires weighing scales.

A number of websites have sprung up recently claiming to calculate a user's personal carbon footprint. For example, AMEE[1] calls itself 'the world's energy meter' while Encraft[2] 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.

[1] `http://www.amee.cc`
[2] `http://www.encraft.co.uk`

2 Potential benefits and drawbacks

PEM data will enable us to identify areas for optimising our consumption of resources. It may make metering and billing by energy companies '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 PEM 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 facility. Many energy reducing measures also bring monetary savings, but there is insufficient information available to let consumers realise their best courses of action. I hope that a PEM could help its users make more informed decisions.

There is much discussion in the press of 'carbon footprinting', but 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. More careful apportionment can be applied to data regardless of how it is collected and offers more meaningful results.

3 Promoting wide-scale adoption

Ubiquitous computing has developed along slightly different lines from those envisaged when the term was coined over two decades ago [5]. Mobile phones are now an almost ubiquitous computing platform, and they are increasingly becoming powerful sensor platforms as well. I imagine that future generations of mobile phones might contain a PEM as an integral part.

Social networks provide an ideal forum for users to share consumption patterns and reduction strategies, and I have begun exploring the significant overlap between the fields of sentient computing and social networking [6].

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[3] 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[4] 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. Dopplr[5] lets its users share estimates of their carbon emissions through travel, and DIY KYOTO[6] has a community site based on its electricity monitoring device, al-

[3] http://www.nikeplus.com
[4] http://www.dailymile.com
[5] http://www.dopplr.com
[6] http://community.diykyoto.com/

lowing 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, highlighting the global significance of seemingly small scale actions.

I hope that the power of the social networking phenomenon might be brought to bear in a similar way to drive adoption of the PEM and provide impetus and support for changing lifestyles.

4 Strategy

Fig. 1. Estimated energy consumption of a "typical moderately-affluent person"

It will probably be infeasible to measure and apportion all of the many ways in which we consume energy. I am therefore beginning my investigation around the most significant areas of consumption as estimated by MacKay [7] (Figure 1). My research into the PEM builds on existing efforts for environmental footprinting by considering the technology necessary to divide the estimated footprint for a building or organisation amongst individuals.

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. 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.

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. 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.

Heating, cooling Automated recording of the electric and heating costs for a building is increasingly common. However, there are many plausible schemes for apportioning this cost to building users. For example, costs could be divided among those entitled to use the facilities, or split based on physical occupancy time. The PEM must be able to accept building specific allocation policies and must be able to access the contextual information required to implement them.

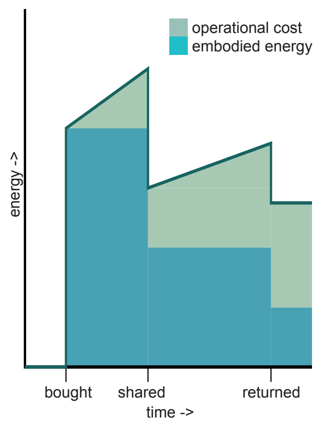

Fig. 2. Total energy cost of ownership

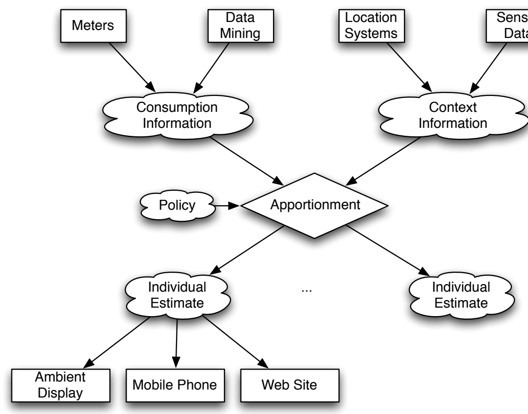

Fig. 3. Example personal energy meter architecture

Stuff The energy consumption trace for an appliance (Figure 2) 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. 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 thus take a share of its energy cost.

5 Architecture

The design and implementation of a PEM embodies many challenges. Input data must come from a wide range of meters, sensors and location systems distributed globally; apportionment 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.

Figure 3 illustrates some of the data flows envisaged in a PEM architecture.

6 Apportionment

Imagine a coffee machine shared between members of a team. It seems logical that a person ordering a drink should be allocated the electrical energy consumed in making it—but he should also be responsible for a portion of the energy needed to keep the water hot throughout the day, and a portion of the energy used processing and transporting the coffee beans from where they were grown.

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 [8]. 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.

Context awareness offers the potential to change this balance and apportion energy costs to those who cause them to be incurred: the person standing at the cooker should be responsible for the gas it consumes during that period (though these costs may be reallocated if the meal is shared), and the cost of the electricity required by a television should be split between all those sitting in front of it.

Apportionment is one of the greatest challenges, but also promises some of the most significant benefits of a personal energy meter.

7 Specifying policy

Apportionment policies may vary not just in different scenarios but even from institution to institution, building to building and object to object. For example, one company might decide to apportion heating costs for a building in proportion to the amount of office space allocated exclusively to each employee, while another might split it based on the amount of time spent in the building. We need a language for specifying these policies in terms of the contextual information that drives them. Model checking techniques might be applicable to ensure that desirable global invariants hold: for example, that the sum of all the energy apportioned to individuals is equal to the total energy measured.

8 Context awareness

Contextual information will provide 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 [9]. Indoor location systems have been the subject of much research over the past two decades [10], 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 will be low cost, low infrastructure location systems that can be deployed on a truly global scale.

8.1 Low infrastructure location systems

One approach to reducing installation costs is to replace the expensive custom hardware employed by conventional solutions with commodity devices. For example, optical motion capture and tracking systems provide very high levels of accuracy, but generally use bespoke cameras and cost thousands of pounds. I have demonstrated a method for using Nintendo Wii controllers as a stereo vision system to perform 3D tracking in real time [11]. In this case, off-the-shelf consumer hardware allows a wireless, portable tracker to be created that obtains accurate results for a fraction of the cost of conventional setups. Consequently, tracking becomes viable in situations where cost or space were previously prohibitive.

Another parallel strategy involves leveraging existing infrastructure that has been installed in the building for a different purpose and exploiting its characteristics to obtain contextual information. For example, activity recognition has been explored using a small number of microphone-based sensors at critical locations in a home's water distribution infrastructure [12], and the PowerLine Positioning system claims sub-room-level localisation using plugin modules that inject a low frequency, attenuated signal throughout the electrical system of the home and active receivers that listen for them [13].

Radio-based systems show significant promise in this area, and functional indoor location systems have been constructed using WiFi [14], GSM [15] and Bluetooth [16, 17]. Bluetooth-based systems are particularly attractive as they have low power requirements and almost everyone already carries a mobile phone and has a computer on his desk; however, location systems to date have relied on the inquiry mode to constantly scan for devices. This process is very slow, can be a security and privacy risk and is not supported on many new handsets. I am studying the use of low-level Bluetooth connections to track mobile devices within a field of fixed base stations. The connections avoid the use of Bluetooth discoverability altogether and are centrally co-ordinated. I have investigated the properties of the low-level connections both theoretically and in practice, and am constructing a building-wide tracking system based on this technique.

8.2 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 very energy efficient ways and is operational only when delivering a service for some real end-use [18].

Recently, Google engineers have argued the case for energy proportional computing in large server farms [19] and corporations are beginning to wake up to the potential cost savings of reducing their energy consumption; however, very little consideration has been given to the cost of location systems. For example, the Bat system relies on a network of ultrasound receivers installed in the ceiling [20]. These receivers are permanently on, drawing approximately 25W per room. In the past, most work relating to energy consumption of location systems has focussed on extending the battery life of mobile devices rather than reducing the global cost.

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 localization network [21] and Nishida et al. investigated the number of receivers required in ultrasound-based systems [22]. These are generally aimed at simplifying administration rather than reducing energy consumption.

A global indoor location system should strive for energy proportionality and be truly *agile*: it must scale its performance—and power requirements—in order to meet the demands placed on it at any given time. When there are few clients, or only rough estimates of position are required, sensors should be switched off; when there is a need for more accurate answers they can be brought back into service.

In order to achieve this, the systems will have to be *dependable* [23] so that their performance can be measured against a required standard. Different applications place different demands on the system: for example, virtual buttons as implemented by the Bat system require both a high resolution and a high degree of confidence in the position estimate, as several buttons may be located very close together, while telephone forwarding would only require room level answers and may be able to reduce the update rate dramatically for stationary users.

9 Conclusion

A personal energy meter that provided live information on consumption apportioned to individuals would represent a very significant step forwards from the current best possible situation of a static, approximate and time consuming audit of a building or organisation. There is a number of challenges to be addressed, including a language for specifying apportionment policies and systems to provide the necessary contextual information. I hope to make a valuable contribution to the state of the art in sentient computing as well as raising awareness of the global energy issue.

Acknowledgements

I wish to thank Alastair Beresford, George Coulouris, Robert Harle, Joseph Newman and Andrew Rice for their continued help and advice and my supervisor, Andy Hopper, for his insight and support.

References

1. Kitzes, J., Wackernagel, M., Loh, J., Peller, A., Goldfinger, S., Cheng, D., Tea, K.: Shrink and share: humanity's present and future ecological footprint. Phil. Trans. R. Soc. B **363**(1491) (Jul 2007) 467–475
2. Hopper, A., Rice, A.: Computing for the future of the planet. Phil. Trans. R. Soc. A **366**(1881) (Oct 2008) 3685–3697
3. Kurp, P.: Green computing. Communications of the ACM **51**(10) (Oct 2008)
4. Jochem, E.: Steps towards a sustainable development. A White Book for R&D of energy-efficient technologies
5. Bell, G., Dourish, P.: Yesterday's tomorrows: notes on ubiquitous computing's dominant vision. Pers Ubiquit Comput **11**(2) (Jan 2007) 133–143
6. Hay, S., Newman, J., Rice, A.: Sentient computing meets social networking. W3C Workshop on the Future of Social Networking
7. MacKay, D.J.C.: Sustainable energy—without the hot air. UIT, Cambridge
8. Gneezy, U., Haruvy, E., Yafe, H.: The inefficiency of splitting the bill. The Economic Journal **114**(495) (2004) 265–280
9. Dey, A.K., Abowd, G.D.: Towards a better understanding of context and context-awareness. The What, Who, Where, When, and How of Context-Awareness
10. Hightower, J., Borriello, G.: Location systems for ubiquitous computing. Computer (Jan 2001)
11. Hay, S., Newman, J., Harle, R.: Optical tracking using commodity hardware. ISMAR 08
12. Fogarty, J., Au, C., Hudson, S.E.: Sensing from the basement: A feasibility study of unobtrusive and low-cost home activity recognition. UIST '06
13. Patel, S.N., Truong, K.N., Abowd, G.D.: PowerLine Positioning: A practical sub-room-level indoor location system for domestic use. UbiComp '06
14. Bahl, P., Padmanabhan, V.: RADAR: an in-building RF-based user location and tracking system. INFOCOM '00
15. Otsason, V., Varshavsky, A., LaMarca, A., de Lara, E.: Accurate GSM indoor localization. UbiComp '05
16. Anastasi, G., Bandelloni, R., Conti, M., Delmastro, F., Gregori, E., Mainetto, G.: Experimenting an indoor Bluetooth-based positioning service. ICDS '07 Workshops
17. Hallberg, J., Nilsson, M., Synnes, K.: Positioning with Bluetooth. ICT '03: Proceedings of the 10th International Conference on Telecommunications **2**(23) 954–958
18. Hopper, A.: Opinion first person—if you ask me—computing the planet's future. Engineering & Technology **2**(7) (Jul 2007) 24 – 24
19. Barroso, L.A., Hölzle, U.: The case for energy-proportional computing. Computer **40**(12) (2007) 33–37
20. Addlesee, M., Curwen, R., Hodges, S., Newman, J.F., Steggles, P., Ward, A., Hopper, A.: Implementing a sentient computing system. IEEE Computer **34**(8) (Sep 2001) 50–56
21. Jevring, M., de Groote, R., Hesselman, C.: Dynamic optimization of Bluetooth networks for indoor localization. AASN '08
22. Nishida, Y., Aizawa, H., Hori, T., Hoffman, N.H., Kanade, T., Kakikura, M.: 3D ultrasonic tagging system for observing human activity. IROS '03
23. Rice, A.C.: Dependable systems for sentient computing. Computer Laboratory Technical Reports (2007)