

The Potential for Location-Aware Power Management

R. K. Harle
Computer Laboratory
University of Cambridge, UK
rkh23@cam.ac.uk

A. Hopper
Computer Laboratory
University of Cambridge, UK
ah12@cam.ac.uk

ABSTRACT

This paper explores the use of location-awareness to dynamically optimise the energy consumption of an office. It makes use of high-accuracy location data collected over 60 days randomly selected from a year in a commercial environment to evaluate the potential for energy savings and to motivate techniques that might be used.

The results suggest that the energy expended on lighting and fast-response systems could have been cut by 50%; that 75.8% of the average user's working day was spent in their office; and that around 140Wh per PC per day could have been saved, compared to a policy that had machines on for the entirety of the working day. We also find inconsistent office usage that would make optimising slow response systems much harder.

Author Keywords

Energy Efficiency, Location-aware, Context-aware, Power Management

ACM Classification Keywords

I.6.3 [Simulation and Modeling]: Applications; H.5.2 [User Interfaces]: Evaluation/methodology;

INTRODUCTION

The explosion in consumer electronics has carried with it a price in terms of increased energy demands in all aspects of our lives. The environmental costs of producing the necessary energy are coming under scrutiny, and we are seeing a push toward reductions in power usage. Generally, the energy reduction techniques come at the expense of our convenience or comfort: we are advised to shower for shorter periods; avoid using standby states on electronics; set the thermostat a few degrees cooler. Whilst these approaches undoubtedly have a beneficial effect in reducing energy consumption, researchers have also recognised the potential to use technology to monitor other technology, with sensors providing a view of the world from which the environment might be autonomously optimised. In principle, this could

permit for reductions in energy consumption without reducing convenience and comfort — a ‘have your cake and eat it’ approach.

Implementing a system to dynamically optimise an environment is recognised as a very difficult problem. The most common approach has been to use a variety of simple sensors and apply Bayesian inference techniques to attempt to infer the state of user(s) in the environment and hence adapt the power states of associated devices. For example, [6] looked at desktop PCs and tried to dynamically and probabilistically assign each user a state from the set *{using, about to use, not using}* in order to set the power state of the associated PC appropriately.

In this paper, we contribute to this field by studying the extent to which fine-grained location throughout a building might be used in dynamic power management of devices. To make our observations, we use a large corpus of captured location data from a non-academic environment with 50 rooms, involving around 40 people with various roles, captured during 60 working days randomly selected from a year. The data were not collected specifically for this work, but rather as a corpus of data for future, unforeseen tasks such as this work. Consequently, this analysis is not a detailed evaluation of a location-aware power management system, but an analysis of what such systems might be able to contribute. We analyse the collected location data to form a picture of how people work and what energy savings might reasonably be expected if we were able to prevent device ‘idling’.

In particular, we emphasise that any changes should not frustrate users. As an example, we informally asked a number of computer users in various commercial and non-commercial settings why they did not have their desktop machines automatically power down, suspend-to-disk, or go to sleep after a specified time. The majority cited the frustration of having to wait for the system to reach a usable state as a major contributing factor. We cannot afford, then, to assume that an aggressive power saving policy will necessarily lead to power savings since it may prove very unpopular and be circumvented. Instead workstations in the scheme discussed must be powered up and in a usable state *before* the user is physically upon them.

BACKGROUND

Context-aware power management (CAPM) was introduced by Harris and Cahill [5, 6], who concentrated on using sen-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

UbiComp'08, September 21-24, 2008, Seoul, Korea.

Copyright 2008 ACM 978-1-60558-136-1/08/09...\$5.00.

sors to optimise the power states of desktop PCs in an office environment. They used coarse location in the form of proximity detection using Bluetooth, microphones to detect activity, and webcams and facial tracking to estimate when a display was being observed. They characterised users as HeavyUse ($\geq 85\%$ of their time using their computer when in its vicinity) and LightUse (significantly less time spent using their PC when in its vicinity). They found that predicting user state for HeavyUse was generally successful (within 8% of the ground truth ‘oracle’), whilst achieving useful prediction for the LightUse grouping was essentially too challenging for today’s systems.

Our interpretation of these results is that predicting when a computer will be used by a nearby user is not a realistic goal for today’s technology. The time taken to transition between PC power states is significantly longer than the time taken for a user to shift task from, say, reading a book to typing. Instead we posit that the success of the HeavyUse user analysis is due entirely to the Bluetooth proximity detection — i.e. it is only currently possible to predict device usage by assuming it is in use when a user is within operational range of it. Since HeavyUse users essentially use their PC whenever near it by definition, the resultant predictions were close to the oracle.

Furthermore, even if we were able to accurately predict when a computer was not about to be used, entering a low power state does not support the manner in which we have empirically observed modern PCs to be used. The proliferation of global connectivity means that desktop PCs¹ are often used for monitoring tasks when not in physical use. When at their desks, users want to know when email arrives and when instant messaging sessions are requested.

In this paper, we use these observations to motivate a single assumption: *when a user is within range of a device, the device should be treated as in use*. This means we do not strictly minimise energy usage, but we avoid ever inconveniencing the user, whilst still making significant energy savings. It also permits us to be more reactive in our approach.

Note that some devices are able to change the power state of specific components very quickly. A PC, for example, may power down its display or its hard disk and resume both extremely quickly. In recognition of this, we adapt our definition of ‘in use’ to permit these changes since they cause minimal frustration. In effect, we are merely saying that an ‘in use’ device should behave exactly as is acceptable to the user today — i.e. *we are attempting to apply location-aware power management such that the results are invisible to today’s end users*.

Harris and Cahill’s work is the only one we are aware of that concentrates on the office environment. CAPM in the

¹We assume that a desktop PC is wholly dedicated to providing local computing and general network services are provided by carefully monitored machines in controlled server rooms. This approach is more energy efficient and appropriate for business.

home has, however, received much attention, particularly in the MavHome project [11]. Here, researchers developed an information-theoretic framework for predicting the motions of users in a home environment. Sensor input again concentrated on location, this time using RFID readers and pressure mats. Each location within the home was assigned a unique textual descriptor and any movement was represented by concatenating these descriptors into a string. Textual compression schemes were used to highlight patterns in the strings and hence assist in predicting where any user was heading to allow for optimisation. The framework was evaluated only in simulation for a small house, which makes it difficult to know whether the prediction scheme could function well enough to avoid frustration.

AI techniques have also been used to analyse and predict building usage. Research such as the Neural Network House [9] and by Mayhofer et al. [8] have assumed sensor-laden homes that can make semi-intelligent decisions in order to save energy (e.g. if no-one is home at 11pm, no-one is likely to be staying there tonight so switch to a low power mode). Again, though, these tend to be applied to small residential dwellings and evaluated through simulation. We view this work as complementary rather than competitive.

We note that any probabilistic approach (such as used in MavHome or by Harris and Cahill) will occasionally cause a machine to be in an unusable state when a user wishes to use it and therefore will cause frustration sooner or later. We hope to demonstrate here that for some energy sinks this is inevitable and the probabilistic approach is appropriate, whilst for other devices we can implement schemes that will provide no frustration at the cost of *slightly* greater energy usage. Again, this highlights the trade-off between convenience and power minimisation, but we feel that the goal should be to minimise power under the constraint of maximal convenience. If this is not the case, schemes will be circumvented or simply not implemented.

Since we make use of fine-grained tracking systems, it is also worth providing a brief review of the state of the art in order to justify their consideration. Demonstrated in-building systems have been based on many different approaches, each with important characteristics. Infra-red badges [14] provide reliable room-level location with low infrastructure costs; WiFi radio systems [3, 15] have very low infrastructure and deployment costs, but crucially they are unable to reliably associate a user with a room because the signals penetrate walls; ultrasound systems [2, 7] have very high accuracy (of the order of centimetres) but equally high infrastructure costs; UWB radio systems offer a good compromise with sub-metre tracking and intermediate deployment costs [13]. UWB systems in particular are becoming popular in commercial and industrial settings.

UNDERSTANDING WHERE THE ENERGY GOES

In order to understand the true effect any optimisation might have, it is important to consider where energy is typically used. Naturally, the power budget for a given building is dependent on a large number of factors, from size to func-

Source	Residential	Commercial
Space Heating	30.7%	14.2%
Space Cooling	12.3%	13.1%
Water Heating	12.2%	6.8%
Ventilation	n/a	6.0%
Lighting	11.0%	25.5%
Refrigeration	7.5%	4.1%
Electronics	7.4%	6.3%
Wet Clean	4.8%	n/a
Cooking	4.5%	2.0%
Computers	1.1%	3.2%
Misc	8.5%	18.7%
Total Sector Expenditure in Petawatt hours(PWh)	6.38154	5.2447

Table 1. Typical building power expenditure (Source: US DoE [10]).

tion to employee count to the current season. Table 1 shows the primary sources of energy expenditure are environmental temperature control and lighting, but many express surprise that the consumption attributed to computers is so low in comparison. This is partly due to the fact that computer equipment receives a bad press: it is a target because it represents something that is visibly consuming power when idle, and can therefore be ‘easily’ optimised. These data motivated us to expand our interest from the desktop PC to look at whether savings might also be possible for other major sources of energy consumption. Since our location data is from an office environment, this is our target herein.

Some of the devices listed in Table 1 are not optimisable given user locations. There is potential to dynamically optimise heating, cooling, lighting, and electrical devices, but those devices that either demand human presence or are indifferent to it (such as refrigeration, cooking and cleaning) have little chance of being dynamically optimised using a CAPM system and hence receive no further consideration here.

GROUNDING OUR FINDINGS

Providing meaningful answers to our research questions is challenging at best. To ground our findings we avoid simulating building use, and opt instead to analyse the *actual* usage of a specific building with an installed personnel tracking system. The building used was the former AT&T Research building in Cambridge, UK, which featured an ultrasonic location system that provided three-dimensional tracking accuracy of approximately 3cm 95% of the time [2]. The building had three floors and 50 rooms, with a mix of users from management, research, and administration. We classified each room as an office, a corridor, or a communal room (a kitchen, a meeting room, etc). Under this classification, the building had 35 offices, 6 corridor spaces and 9 communal rooms.

Over the course of a year, researchers selected 60 days during which the location events generated by the tracking system were recorded between the hours of 0630 and 1930. The resultant logs provide a detailed picture of where building users went, and are used throughout this paper as a representative view of the use of an office environment. Clearly, no two workplaces are the same, so it is difficult to con-




Figure 1. Positions recorded for a particular user over four hours on a particular day. The data were recorded across three floors, each shown separately.

fidently generalise the results, but qualitatively we observe trends typical of other workspaces and we feel that the results derived from them are more reliable than the alternative of simulation.

It is, however, important to understand that the location logs cannot be guaranteed as complete or totally accurate. Building users were not required to submit to being tracked, nor were they required to ensure that they wore their tracker at all times. However, there were incentives to be tracked and there is no evidence to suggest that the data are not a good representation of how the office space was used. Figure 1 illustrates the general quality of the data by plotting the logged path for a particular user during the course of a day.

CLASSIFICATION OF ENERGY SINKS

When considering typical energy sinks and how to optimise them, we found it useful to categorise them by their ‘resume’ times from a reduced-power state. Broadly speaking, we consider three categories:

1. Instantaneous resume. To the human eye, these devices switch power state instantaneously. Lighting and electronic displays are good examples.
2. Sluggish resume. These devices have non-instantaneous resume times, but still resume within the order of a minute. A typical example is a desktop PC.
3. Slow resume. These devices are associated with slow warm-up periods often of the order of fractions of hours. Heating and cooling systems are good examples.

OPTIMISING INSTANTANEOUS RESUME

Many buildings are already equipped with coarse location-based lighting control in the form of wall switches with embedded infra-red detectors. When these sense the presence of a heat source (presumed to be human), the local lighting is switched on. Similarly, the absence of a heat signature triggers a countdown to powering off the lights. Naturally this approach is more energy efficient and convenient than a physical light switch, but in our experience the approach suffers from a number of drawbacks:

- The scheme is purely reactive. To be recognised one must




Figure 2. Using spatial zones as triggers. Rooms A and B are connected as shown. Each has two associated spatial zones (shown hatched), one for ingress and one for egress. The ingress zone is chosen to spill into the adjacent room to allow warning of imminent likely entry.

be in the room, which necessitates entering a dark space initially.

- Even if detected quickly, the lighting system may not respond instantly. This is especially true of power-saving bulbs, which tend to require a short warm-up period.
- It is hard to completely cover an entire space in such a way that users cannot inadvertently hide from the sensor. This can cause the lighting to turn off at inappropriate moments.

These problems are barely tolerable for a lighting system (in our laboratory, they are often circumvented to avoid being plunged into darkness at inopportune moments). The lack of any notion of identity is a further issue for more personal devices — every PC display in the room or area should not spring to life when a cleaner enters. There is, then, good motivation for using a more fine-grained tracking system that can identify those it tracks. Instantaneous resumes lend themselves to predominantly reactive schemes, although a degree of pro-action may be desirable.

To handle reactive events we use the notion of spatial zones, first applied by the Sentient Computing project developed by the University of Cambridge and AT&T Research [2]. The principle is simple — define 2D spatial zones (essentially polygons overlaid on a map of the area) that trigger system events whenever a user enters or exits them [4]. Fine-grained tracking is important here for two reasons: firstly, it is usually associated with faster location update rates and hence better estimation of the time at which zones become ‘active’; secondly, it provides a more reliable estimate of proximity. Coarse proximity-based location such as that employed by Harris and Cahill is usually unable to associate a user with a room and thus resumes could be mistakenly triggered when the user is actually in an adjacent room.

When dealing with lighting, we use the room outlines as the relevant spatial zones, and take advantage of the faster update rates associated with fine-grained tracking to minimise the delay in turning on lights as a user enters. However, in some cases, any delay at all may be unacceptable, in which case we extend the notion of spatial zones to incorporate a separate ingress and egress zone, where the latter is the true room zone and a subset of the former. The ingress zone is an extension of the egress zone at the doorways to create small ‘spill-over’ areas which will be triggered just *before* the user enters, allowing for pro-active resumption of lighting in the




Figure 3. Captured location sightings in a corridor show the tendency for users to adopt the corridor centre line when traversing buildings.

room if it is currently off (Figure 2).

Provided that a given user is reliably tracked, this scheme *guarantees* that the lights will be on when they enter the room, achieving both convenience and invisibility. However, the ingress adaptation introduces the possibility that a resume will be falsely triggered if, for example, a user walked through the ingress zone without subsequently entering the attached room. This would trigger the lights to be on unnecessarily which would be an unwarranted energy consumption. When we reviewed the captured data, we found that this may be less of an issue than first envisaged. Figure 3 shows all captured locations within a particular section of a corridor over the course of a day, and clearly demonstrates the natural tendency for humans to maximise their distance from walls unless diverting to a destination (doorways to attached rooms are shown indented along the corridor outline). To quantitatively assess this, we chose a commonly-occupied room and looked at the average number of false resumes per day that would have occurred for a range of ingress extension sizes. The results in Figure 4 suggest an extension of 0.5m perpendicular to the door would be appropriate for the environment studied. Our empirical results from location-aware systems suggest that the wider the corridors, the greater the extension can be, giving more time for pro-action if desired.

The question of how much of an energy saving this approach provides translates roughly to asking how often a room is occupied. For illustration, we assume three lighting schemes:

1. All lights are on 24 hours a day.
2. Lights are switched on by the first office owner² to enter in the morning and off by the last to leave.
3. Room lights are automatically switched off when the room is vacant using location-awareness.

We filtered the raw location logs using point-in-polygon tests with the room outlines in order to measure the length of time each room was occupied in each day. Figure 5(a) shows the average daily illumination time for each room when using lighting schemes 2 and 3. In recognition of the fact that not all rooms were the same size, we reproduce Figure 5(a) as Figure 5(b), which weights the data according to how many bulbs would be needed to light the room, under the empirical observation that one bulb was required for every 10m².

²We ignore any entry by non-office owners since these were typically caused by e.g. administration putting a folder on a desk early in the morning.




Figure 4. Extending ingress spatial zones and its effect on the number of mistaken resume events.




Figure 5. The average illumination time for each lighting scheme per room. The rooms are presented in order of type: offices (section I), communal (section II), corridors (section III).

In order to generalise our results, we have computed the total daily illumination time for the building in terms of both room-hours³ and our estimated bulb-hours. The results for the different classifications are given in Table 2. Empirically, we note that the majority of non-residential buildings operate on a policy similar to that of scheme 2, so these results suggest the energy requirements for building lighting could be halved by using a reactive location-aware power management scheme.

There is, however, a caveat to this result. It is clear from Figure 5(a) that the biggest advantage to a reactive scheme comes from application to the corridors, but the analysis fails to capture that the small overall occupancy time of corridors is itself the sum of many, much shorter occupancy periods created by user movement. Unfortunately, rapid switching of today’s fluorescent bulbs causes a deterioration in bulb lifetime which may have its own environmental impact [12].⁴ Enforcing a policy that ensures only one state change per assigned period (five minutes is identified as the break-even for fluorescent bulbs in [5]) is extremely difficult for both reactive and proactive schemes, and it is inevitable that there will need to be rapid switching on occasion. We analysed the lengths of the periods of vacancy for each room to discover how likely it is that a given vacancy will last longer than a given time. The results for the different types of room are shown in Figure 6. We observe that the probability a room will remain vacant for at least a 5-minute break-even period is 0.26 (corridors); 0.42 (communal areas); 0.56 (offices); and 0.37 overall. This analysis does not, however, capture the effect of daily routines which may make the probability distribution time-dependent. If, for example, a given office goes to lunch at 12pm each day, it is very likely that any vacancy starting at 12pm will last longer than 5 minutes. We study this in more detail when considering slow resuming systems. In general, for a 5 minute break-even period, we advocate a scheme that reactively turns office and communal lights on for a minimum of five minutes before returning to location-aware management, and which leaves corridor lighting on whilst the building floor is in use. We also note that there is a general trend to replace fluorescent bulbs with LED ‘bulbs’, for which rapid switching is not an issue.

Note that we have assumed lighting is per-room, rather than per-personal area. This is of course the norm today (one lightswitch per room), but we note that location-awareness offers the opportunity to have more specific per-person spotlighting. Similarly, we have noted that some instantaneous resume devices are typically associated with a particular user rather than a particular room (e.g. a display or a telephone). It may be the case that these devices power-down automatically when not in use (computer displays for example), but some electronics provide output on an assumption of the user being there (electronic signs, telephones) and so it is in-

³A room-hour is defined here as an hour of illumination for a specific room.

⁴It is not clear whether the necessary increase in bulb production would outweigh the benefit of reduced energy consumption, but it certainly creates more work for building maintenance departments!




Figure 6. The probability that a newly vacant room will remain vacant for at least a given time.




Figure 7. Proportions of time users spent in their respective offices.

structive to examine whether there is value in location-aware power management of such user-specific devices.

In these instances, we want the spatial zones to be linked to specific users, and these considerations lead us to examine how often users are in their assigned offices. Figure 7 illustrates the percentage of the working day that users spent in their office. From these data, an average worker spent $75.8 \pm 13.5\%$ of their time in their office, showing that, in principle, devices could be switched off for almost one quarter of the working day.

OPTIMISING SLUGGISH RESUME

For this class of device, we consider the desktop PC as an illustrative example. Here, we use the term ‘PC’ to refer to the computer itself rather than its peripherals such as the attached display, since these typically power down and resume near-instantly. Computers typically feature a variety of power saving states that are either transitioned autonomously (processor speed, hard disk spin state) or manually (suspend, power off, etc). In most scenarios, the power savings caused by the former can be significantly outweighed by the latter.

The power saving in most modern day PCs is based around the Advanced Configuration and Power Interface (ACPI) industry standard [1]. This specification defines four global states for a machine; G0 (on or working), G1 (sleeping), G2

Scheme	Average daily illumination			
	Office (Rh/Bh)	Communal (Rh/Bh)	Corridors (Rh/Bh)	All (Rh/Bh)
1	437.5 / 1344	216 / 624	144 / 456	1200 / 2352
2	89.5 / 139.7	74.5 / 200.2	61.5 / 196.6	219.5 / 512.7
3	56.1 / 93.7	30.5 / 115.7	9.6 / 34.1	91.8 / 226.0

Table 2. Average illumination times in terms of room-hours (Rh) and estimated bulb-hours (Bh).

(soft-off), and G3 (hard-off). The G1 state is further divided into ‘S’ states, the most important of which are S3 (‘suspend-to-RAM’, ‘sleep’, or ‘standby’) and S4 (‘suspend-to-disk’, or ‘hibernate’). We model a desktop PC as a machine that operates in states G0, S3, S4, G2, or G3, as per the state diagram of Figure 8. This Figure is annotated with experimentally measured power draws from a Dell Optiplex 745MT E4600 desktop machine running Microsoft Windows XP that implements a modern ACPI interface and enters a true ACPI S3 suspend-to-RAM state⁵. The quoted power values were measured at the mains socket to which the PC was attached.

The, perhaps initially surprising, result from measuring the power draws is that the power difference between the S3 and S4 states is negligible. There is, therefore, a strong argument not to use the S4 state on modern desktops. However, it is important to remember that a machine in S4 draws the same power as in G2 and can be instantly switched to G3 (with its zero power draw) without loss of software state. A networked mains adapter (using a protocol such as X10) can be envisaged that would cut power automatically to a hibernating machine. It is therefore reasonable to proceed on the assumption that S4 draws the same power as G3 (0W).

The natural way to handle a sluggish power state response is to make use of the extended ingress spatial zones, as used in the analysis of instantaneous resume devices above. We can envisage extending the ingress region into surrounding areas until its boundary represents the distance from the room which we expect a user to traverse in the time taken to resume. The logical approach is to assume that users move at their maximum speed to ensure there will always be time to resume. Figure 9 shows the frequency of speeds observed for three random users from a single day’s worth of location data, suggesting a maximum speed of around 2ms^{-1} . This corresponds to a zone-doorway distance of at least 10m for the transition to and from S3, and a huge 120m+ for the S4 transition.⁶

The power scheme naturally becomes the implementation of the following algorithm:

```
if Owner within 10m then
```

⁵In the past, suspend-to-RAM has tended to draw similar power to the idling state, and its usage has thus provided little true benefit. New machines do not suffer from this issue.

⁶The S4 transition time is bounded by the time it takes to transfer the memory image residing on disk into live memory. Modern desktop drives can transfer at average rates around 60MB/s, requiring approximately 17s for each 1GB of main memory in the system. The upcoming generation of Solid State Drives (SSDs) are set to increase the read speed by almost an order of magnitude, however, so S4 to G0 transition times could be dramatically reduced in the medium-term.

```
    Enter G0
else if Owner within 120m then
    Enter S3
else
    Enter S4
end if
```

For many PCs, use of the S4 state is likely only to come when the user has left the building (easily determined with a location-aware system) *or* when the system is able to confidently predict their absence for longer than a previously determined break-even period. Realistically, the best source of such predictions will come from mining shared calendars, which are used extensively in today’s business arena. Knowing, for example, that Alice has a one-hour meeting tabled and has entered the assigned meeting room should be sufficient to confidently predict her absence for at least 45 minutes. Since, however, we are entering the realm of prediction, there will inevitably be mistakes and the invisibility goal will be compromised. Therefore, avoiding the use of S4 when the user is within the building seems to be justified, especially since the difference in power draw between S3 and S4 is almost negligible.

Returning briefly to Figure 7 and the derived result that the average user spent 75.8% of their working day in their office, we see that an average PC can be in S3 for one quarter of the working day, G0 for the rest, and use S4 outside the working day and still achieve the invisibility goal. We have found unattributed quotes that estimate 50% of UK businesses operate with their computers on 24/7, whilst the majority of the remainder operate a ‘switch on in the morning, off in the evening’ policy. If we assume a 5-day working week and an 8-hour working day, with an average G0 power consumption of 100W, an S3 consumption of 3W and an S4 consumption of 0W, the three schemes (24/7, on for working day, location-aware) would consume an average of 2,400Wh, 571Wh, and 437.3Wh per day, per PC, respectively. Location-awareness could therefore provide energy savings of up to 83%.

OPTIMISING SLOW RESUME

Our illustrative example for slow resuming systems is environmental temperature control. Heating or cooling a room is a power-hungry operation that consumes almost one third of the energy used in buildings. The response time from starting a heater or cooler to achieving the desired temperature is dependent on a wide range of complex factors, but is almost always a lengthy process in comparison to other devices. Because of this, reactive power schemes are all but ruled out. The key goal is to predict when a user will next use a given room in sufficient time to allow appropriate heating




Figure 8. State diagram for modelled PC with annotated power draws and transition times where appropriate.




Figure 9. Frequency of speeds observed over a 24 hour period for three users.

or cooling of it. This type of pro-action suits the probabilistic schemes such as those mentioned in the background section above, the techniques for which are well established. Our only contribution to that work is to note that no mention is generally given to the use of shared calendars, which would seem to be a very potent source of a-priori information for the probabilistic frameworks in an office setting.

We note that the degree of success for these probabilistic methods is dependent on how predictable people are in their daily lives. We therefore take the opportunity to see how consistent users in our captured logs were in their use of the building. To study this, we sample the occupancy by anyone $O_{r,d}(t) \in \{0, 1\}$, for each room, r , on each day, d , at discrete one-minute time intervals, $t = t_n$. In order to discover how consistent each period of occupancy was, we further compute a histogram of the resultant occupancy traces, $H_r(t_n) = \sum_{j=\text{days}} O_{r,j}(t_n)$. Figure 10 depicts each H_r as a row in a colour-coded image. We observe that rooms had characteristic start and end times for their usage, which differed from each other by as much as 2.5 hours. Most offices show two consistently occupied periods separated by a mid-day period of vacancy, which we attribute to lunch. Communal rooms received either sporadic, unpredictable use (meeting rooms) or highly consistent use (e.g. a hardware laboratory). The main corridors saw a roughly uniform usage throughout the day, whilst minor corridors had very unpredictable usage.

A typical building heating policy today is to control the temperature of the building between the hours of potential use. For many buildings this corresponds to setting a timer switch

to control temperature between around 0600 and 1930. Figure 10 runs over the same period and it is clear that we can use the observations just made to improve this policy, assuming room temperature can be controlled on a per-room basis:

- We can predict when a user is likely to arrive in the morning such that their office is correctly heated. We can also predict when the first user of the day will arrive to heat the corridors and communal spaces.
- We can observe when an office goes to lunch and predict the length of the vacancy based on historical trends.
- We can observe when users of an office leave for the day and switch off temperature control in that office.

From Figure 10 we observe that no-one was seen before 0700 on any day in any room. Therefore we compare to a best-case timer policy that runs from 0700 to 1930. For just the 30 offices in use, this corresponds to an operational time of 375 room-hours. From Figure 10 we find that this heating policy would control the temperature for 46 room-hours in the morning, 24 room-hours in the evening, and 6.5 room-hours at lunch time when the associated rooms are very consistently empty. i.e. for the offices alone, 20% of the fully operational time is spent heating or cooling rooms that are very consistently empty. This, or course, does not necessarily translate to a 20% energy waste since heating and cooling are very complex processes. Nonetheless, a heating system can operate at a lower power for the predictable vacancy periods, keeping the rooms at a habitable-but-non-optimal temperature in case an unexpected arrival occurs, a policy that

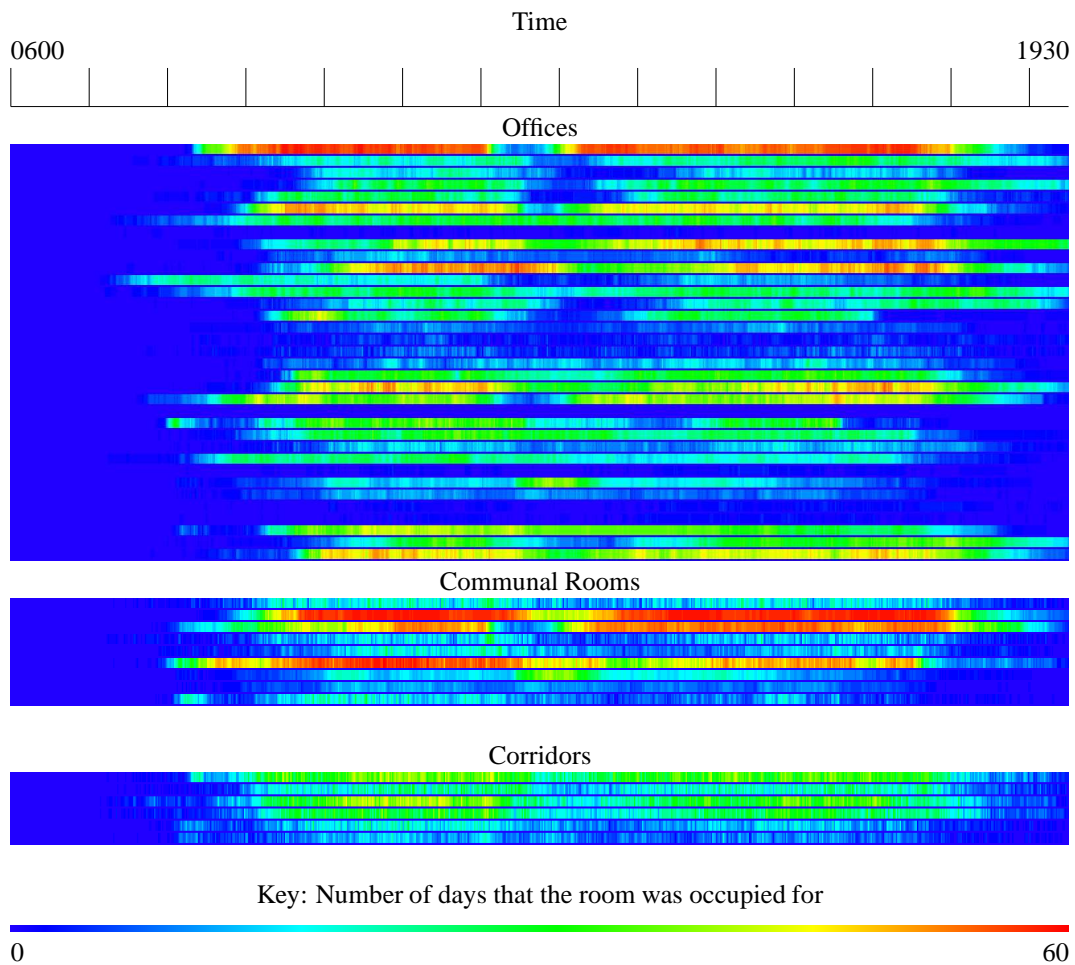


Figure 10. Visualisation of the histograms of room usage across the logged days. The picture shows rooms vertically and time horizontally and colours each 1-minute period of the day according to the number of days that the room was occupied for within that period. With 60 days of logs, this number can be no greater than 60.

would still provide an important energy reduction.

LOCATION SYSTEM POWER DRAW

Throughout this discussion, we have ignored the fact that energy is expended just on determining the location of the users. For reference, the system used to capture the data used in this paper required approximately 40W per room to run, along with one dedicated computer per building. Its successor (the Ubisense UWB platform) draws approximately 14W per room. Whilst these draws are clearly significant, it is important to recognise that neither system has been optimised to minimise power consumption *in any way*, so we would expect careful engineering to *dramatically* reduce this consumption. We also observe that location-awareness should provide other benefits by enabling general location-aware applications and hence this expenditure is not solely for energy reduction.

We also note that, whilst the fine-grained data were necessary to analyse what is possible, some of the proposed schemes could be implemented using much simpler, lower-grade location sensing, which would have correspondingly

lower power consumption.

CONCLUSIONS AND FURTHER WORK

This work has provided a unique look at how a set of 40 people used the building they worked in, with an emphasis on whether building energy consumption can be reduced, primarily reactively, through dynamic optimisation based on input from a building-wide location system. Previous work has used lower-grade sensors and assumed more pro-active schemes, at the cost of making mistakes. Here, we have emphasised that energy consumption needs to be minimised under the constraint of maximal user comfort (or, equivalently, minimal user disruption) relative to today's working practices. We believe this will encourage adoption of our scheme.

Although the data presented are rich, the analysis ultimately only provides a single data point, since it derives from a single building and a single set of people. Nonetheless, in the absence of location data of a comparable nature, we believe the analysis has provided an interesting and novel insight. To summarise the work:

- We believe CAPM systems today will benefit mostly from location data.
- We believe CAPM systems should be designed to reduce energy consumption under the constraint of having an invisible effect to users. This is to improve user acceptance up of the system.
- Predominantly reactive schemes may be more appropriate than pro-active, probability-based schemes since the latter will not guarantee invisibility of the results.
- Lighting is suited to a reactive scheme using spatial zones, especially where LED lighting is used. We estimate that the energy expended on lighting for the building studied could have been cut by around 50% compared to the common first-in-last-out scheme.
- PCs can be semi-reactively optimised by making use of the ACPI S3 power state. We found that 75% of the average user's working day was spent in the vicinity of their computer, and estimate savings of around 140Wh per PC per day compared to a typical scheme that keeps machines on for the working day (and off otherwise).
- The ACPI S4 option for computers is not attractive during the working day due to the time it takes to wake from this state. It is better to use the sleep or standby state unless the return time of the user is known to high accuracy.
- Temperature control has a long response time to power state changes and demands a predictive approach. We have not found sufficient consistency in user behaviour to suggest that predictive approaches will be particularly accurate in practice.

ACKNOWLEDGMENTS

This work was carried out as part of the Cambridge University Computer Laboratory's *Computing for the Future of the Planet* project. Special thanks go to all those who participated in the gathering of the location data. Thanks also to the paper shepherd, Hao-Hua Chu, and all of the reviewers for helping to improve the work.

REFERENCES

1. ACPI. <http://www.acpi.info>.
2. M. Addlesee, R. Curwen, S. Hodges, J. Newman, P. Steggle, A. Ward, and A. Hopper. Implementing a sentient computing system. *IEEE Computer*, 34(8), August 2001.
3. P. Bahl, V. N. Padmanabhan, and A. Balachandran. Enhancements to the RADAR user location and tracking system. Technical report, Microsoft Research, February 2000.
4. Robert Harle. Spatial indexing for location-aware systems. In *Proceedings of the Fourth Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services, (MobiQuitous 2007)*. IEEE, 2007.
5. Colin Harris and Vinny Cahill. Exploiting user behaviour for context-aware power management. In *Proceedings of the International Conference On Wireless and Mobile Computing, Networking, and Communications*, pages 122–130. IEEE, 2005.
6. Colin Harris and Vinny Cahill. *UbiComp 2007: Ubiquitous Computing*, volume 4717 of *Lecture Notes in Computer Science*, chapter An Empirical Study of the Potential for Context-Aware Power Management. Springer Berlin / Heidelberg, 2007.
7. M. Hazas and A. Ward. A novel broadband ultrasonic location system. In *Proceedings of UbiComp 2002: Fourth International Conference on Ubiquitous Computing*, 2002.
8. R. Mayrhofer, H. Radi, and A. Ferscha. Recognizing and predicting context by learning from user behavior. In *Proceedings of The International Conference On Advances in Mobile Multimedia (MoMM2003)*, 2003.
9. Michael C. Mozer. The neural network house: An environment that adapts to its inhabitants. In *Proceedings of the American Association for Artificial Intelligence Spring Symposium, Menlo Park, CA, 1998*, 1998.
10. US Department of Energy. 2007 buildings energy data book. <http://buildingsdatabook.eren.doe.gov>.
11. A. Roy, S. K. D. Bhaumik, A. Bhattacharya, K. Basu, Cook D, J, and S. K. Das. Location aware resource management in smart homes. In *IEEE International Conference on Pervasive Computing and Communications*. IEEE, 2003.
12. E Tetri. Profitability of switching off fluorescent lamps: TAKE-A-BREAK. In *Proceeding of the 4th European Conference on Energy-Efficient Lighting*, 1997.
13. Ubisense. <http://www.ubisense.net>.
14. Roy Want, Andy Hopper, Veronica Falcao, and Jon Gibbons. The active badge location system. Technical Report 92.1, ORL, 24a Trumpington Street, Cambridge CB2 1QA, 1992.
15. M. Youssef and A Agrawala. The Horus WLAN location determination system. In *Proceedings of the 3rd international conference on Mobile systems, applications, and services (MOBISYS) 2005*, June 2005.