The Responsive Environment

Scott Elrod, Gene Hall, Rick Costanza, Michael Dixon and
Jim des Rivieres

CSL-93-5 June 1993 [P93-00015]

CR Categories and Subject Descriptors: **C.3** [Special-purpose and Application-Based System]

Additional Keywords and Phrases: Ubiquitous computing, building management, energy management, embedded computation

XEROX

Xerox Corporation
Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, California 94304

Abstract

This report describes the Responsive Environment Project, an experimental configuration of intelligent offices being built at Xerox PARC. The system gives users a new level of control over their office environments, and demonstrates novel uses of ubiquitous computing for office climate control and energy management. As background, we describe the results of a detailed, one-day occupancy and energy audit. The audit suggests that substantial energy savings are possible with appropriate intelligent control of individual offices appliances. The basic hardware and software are described in detail, as are several software applications. Rationale used in key design decisions are discussed, with particular emphasis on issues that are of general significance to the design of ubicomp systems. We conclude by describing our efforts to address privacy concerns in relation to the gathering of occupancy data.

Acknowledgments

We are grateful to a number of people who have made substantial contributions to this work. Massimo Cavallarro, Janet Davis, David Finch, Al Kachigian, Phil McCoy, Marv Parrish, and Brian Williams all helped with the installation of the hardware in the model area. We thank Bob Wallace of Nova Systems for his help installing and configuring the BMS, and Control Systems International for helping us develop the BMS interface. Scott Clearwater helped us estimate potential energy savings by developing a mathematical model of the air conditioning system, and Austin Henderson and Jeanette Blomberg collaborated in developing the privacy agreement. Finally, we thank the model area occupants for their cheerful participation.

1. Introduction

The Responsive Environment Project was initiated to explore the possibility of making offices that are more comfortable *and* more energy efficient through appropriate use of ubiquitous computing [1]. Because they share the key requirement of seamless integration into the background of ordinary life, ubiquitous computing and office environmental control are natural partners. In order to explore the use of ubicomp for office control, we have built a flexible hardware and software testbed (the *model area*) consisting of 13 offices and some adjacent public areas. The model area provides computer control and monitoring of temperature, room lighting (both natural and artificial), and office appliances (computer monitors and task lighting) within each office. In conjunction with the pre-existing computational infrastructure of workstations, active badges, and mobile computers, this will allow us to experiment with a wide variety of control strategies, including:

- allowing the occupant to set a desired light level, and then automatically adjusting window blinds and/or overhead light intensity to maintain that level throughout the day;
- selecting different light levels depending on whether the occupant is working at her computer monitor or desk;
- providing real-time feedback on the extent and sources of current energy consumption;
- automatically switching off lights and appliances when the room is unoccupied; and
- reducing heating or air conditioning requirements when available information (e.g. from active badges or calendars) suggests that a room will be unoccupied for an extended period of time.

Through these experiments we hope to learn what types of control strategies users prefer, what kinds of control and feedback are appropriate, and how much energy can be saved by such strategies.

2. Background: Work Practice and Energy Savings

A key motivation for this project is the possibility of significantly reducing energy costs. Through extensive use of a computerized building management system (BMS), Xerox PARC has been able to demonstrate an energy savings of 45%, corresponding to more than $2.00 per square foot per year [4]. (See [2] for an excellent introduction to the burgeoning field of computerized building management systems.) Most of this savings has been achieved by upgrading the principal components of the heating and air conditioning system (fans, pumps, chiller, etc.), and through novel control strategies implemented with the BMS. For the most part, however, the building is still controlled as a single unit. For example, the system attempts to maintain a uniform temperature of 73.5° throughout the building. As described below, many offices are vacant a substantial fraction of the time; this offers the possibility of achieving considerable additional savings through intelligent control of individual office appliances and air delivery.

In order to evaluate the potential savings from occupancy-based control of office air delivery and appliances, a one-day audit was conducted in PARC's Computer Science Laboratory.

During the 14 hour survey, 74 areas were monitored hourly for occupancy as well as for all aspects of energy use. Table 1 summarizes the important features of the study.

Rooms under study:	56 individual offices 3 conference rooms 3 laboratories 9 open areas 3 others
Duration of study:	6:00 am – 9:00 pm
Measured quantities:	Occupancy (yes/no) Overhead lights (on/off) Computer monitors (on/screensaver/off) Task lighting (on/off) Other appliances (on/off) Curtains closed (degree) Sunlight in room (low/medium/high)
Total computer monitors:	103
Total CPUs:	99
Total personnel housed in area:	66
Personnel away on day of survey:	15

Table 1: A one-day energy audit conducted at PARC

Figure 1 shows the various contributions to electrical energy use during the 14 hour period. As revealed in the figure, computer monitors and lights represent a substantial fraction of the energy budget. The dip in energy use at 2:30 occurred during a hiring meeting, when most members of the laboratory were together in one large conference room. The dip resulted from a small energy savings as computer monitors automatically went into screensaver mode, and from some lights and appliances that were intentionally turned off. From the graph, one can discern a substantial rise (10 kilowatts) in energy use during the day, but with a considerable background usage of 27 kilowatts that persists during the night.

In addition to the hourly audit, we collected fine-grained (15 second interval) active badge [5] data for 10 individual offices as shown in Figure 2. As can be seen in the figure, there is wide variation among individual work practices. From the active badge data, we can infer an average workday length of 8.8 hours (measured as the time from first arrival in office to final departure from office). A much lower value of 6.5 hours was deduced for the same 10 individuals from the hourly survey, most likely due to the granularity of the measurement. During the 14 hour period, the average occupancy rate for the 10 individual offices was measured at 20% from active badge data and 31% from the hourly survey. Using the hourly survey results, the occupancy rate for all 56 offices combined was 21% of the 14 hour period.

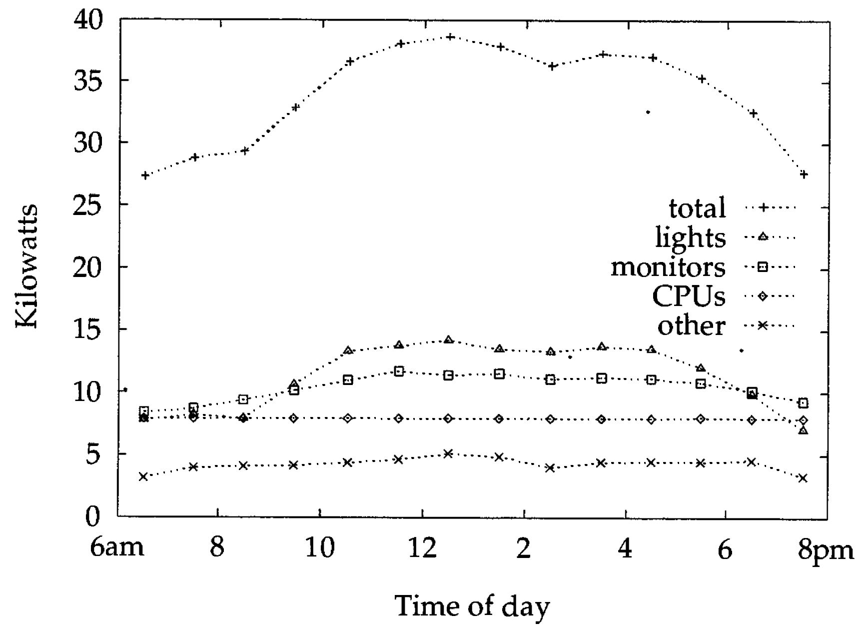

Figure 1: Energy use in the model area during the survey

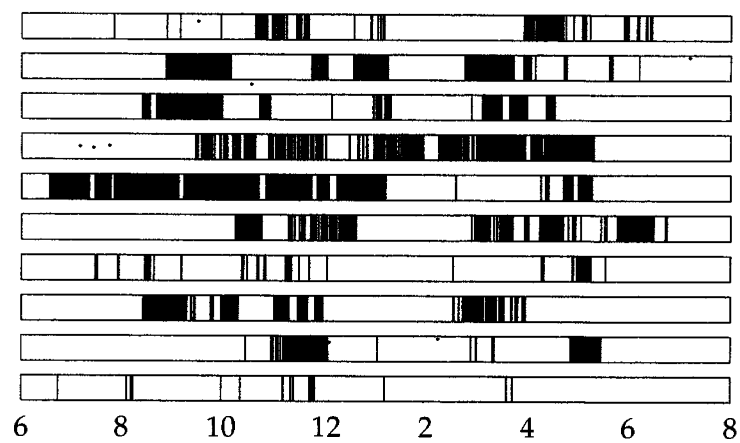

Figure 2: Occupancy data for ten different offices over the course of a single day. Each bar is shaded when the corresponding office is occupied and blank when the office is vacant.

Figure 3 shows the distribution of vacancies for the combined data from 10 active badge wearers. The plot shows the total amount of time spent away from the office (Y axis) for excursions lasting a certain length of time or longer (X axis). The figure shows data for both the single survey day and a 5 day study of the same 10 offices. Using the 5 day average, the data shows that an individual will typically spend 3.5 hours away from the office on excursions lasting 30 minutes or longer.

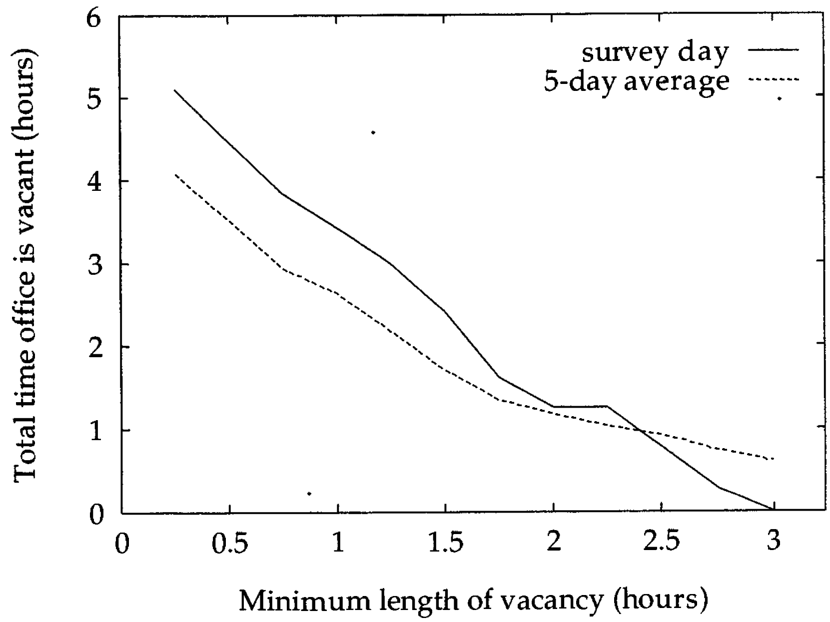

Figure 3: Total time spent away from the office on trips of at least a given length.

Sensible strategies for saving energy will depend in detail on the work practices revealed in Figures 2 and 3. Many factors will be relevant in deciding whether a particular appliance should be turned off during a vacancy of a given length. As an example, certain appliances may exhibit reliability problems if they are power cycled too frequently. In addition, there may be energy costs associated with transients as devices are turned on. In the case of air conditioning, there is a fundamental time constant associated with the thermal inertia of the room. Typically, it is not possible to bring the temperature back to a desired value in less than 30 minutes. For all these reasons, the energy saving strategy must be tailored to each appliance, and will depend upon the detailed occupancy statistics for a given office.

Using the hourly occupancy data, and neglecting the subtleties alluded to above, we can compute the savings that could be achieved simply if lights and monitors were turned off when offices were vacant. The result is shown in Figure 4, along with the total measured energy consumption. On average, 12 kilowatts could be recovered by using occupancy-based control of lights and monitors. For all of Xerox PARC, this extrapolates to an annual savings of $55,000.

In conclusion, there is the potential for substantial savings from intelligent control of appliances and air delivery on an individual office basis. Because of the complexities of the actual systems (latency, reliability, transient start-up costs), appropriate control will require detailed knowledge of occupancy patterns.

3. Hardware

As described above, the model area is a set of 13 offices and an open area that have been outfitted with remote sensors and actuators to control comfort and energy consumption. Currently, the system allows for control of temperature, lighting and desktop appliances.

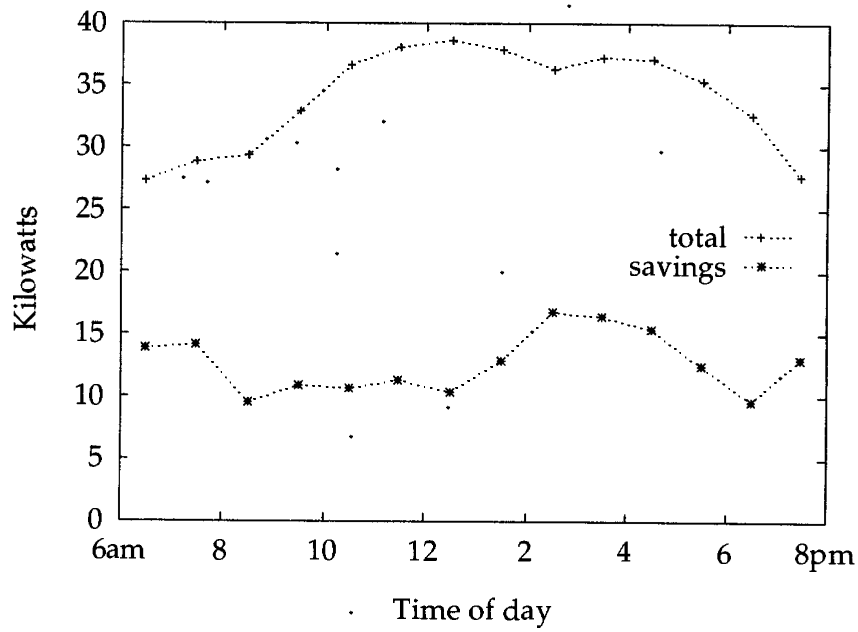

Figure 4: Potentials savings from occupancy-based control of lights and computer monitors

Temperature is controlled by the system in a fairly complex and indirect manner. Each office is equipped with a *variable air volume* (VAV) controller, a vent which can be incrementally opened or closed to admit air from the supply air ducts. The supply air is a mixture of fresh outside air and recycled inside air, heated or cooled as appropriate. Thus opening up the VAV both increases the fresh air supply and moves the office temperature closer to the supply air temperature.

Lighting control includes both artificial and natural light sources. Artificial light is provided primarily by dimmable overhead fluorescent lamps. To improve both comfort and energy efficiency, we are experimenting with high-efficiency natural color lamps and indirect lighting fixtures. We have installed hardware to switch the lights on and off under computer control and to allow continuous dimming down to 20% of full power.

In addition to fluorescent lighting, we plan to control natural light from the windows by using motorized miniblinds. This is important both for comfort (to control glare from computer screens) and to reduce air conditioning demand during the summer.

As pointed out in the discussion of the energy survey, computer monitors consume an appreciable fraction of the total electrical power. We are installing hardware to allow remote control of power to monitors, task lighting, and other AC appliances.

Each office is equipped with commercially available sensors to measure temperature, visible light level and room occupancy. In addition, a number of other sensors have been placed at key locations around the building to allow for precise measurement of external conditions and energy consumption.

Controls: Hardware versus Software?

Determining the balance between hardware and software controls on ubicomp systems represents a key design decision. For greatest flexibility, soft controls are highly desirable. In some cases, however, exclusive use of software control may be incompatible with the key ubicomp goals of being "invisible", and "blending into the woodwork". In the context of the model area, there are several cases where this goal of invisibility has led us to implement hard controls. We believe that these considerations are of general importance in the design of ubiquitous computing systems.

1) The software applications that give users control over their offices will be built on top of a large computational network, which experiences occasional failures. Those features which are crucial to use of the office must have hard controls in addition to any software controls. The lights and computer monitors definitely fall into this category.
2) Users have long-standing habits of use which should be supported in an intelligent office. In particular, people often hit the light switch as they walk into their office. While in some cases users may choose to rely entirely on automatic control of lighting, the controls in offices should also support familiar practice.
3) Up to 100 people change offices each year at PARC. Some means is needed to establish a basic, known level of functionality when a new tenant or temporary visitor arrives. Using a hard control ensures that it can be easily activated by the facilities people responsible for setting up offices.
4) Since the software will be constantly changing, there may be times when it is unstable. In such situations, the office tenant should be able regain control and have the office perform at least as well as before automatic control was installed.
5) People may have privacy concerns about the use of occupancy sensors in offices. A direct physical means of disabling the occupancy sensor within the office might be helpful in giving the occupant a greater sense of control.
6) Users expect certain functions to happen immediately. In our case, software control through the BMS can exhibit latencies of up to 2 seconds. Users would definitely find this delay annoying if it occurred when the light switch was activated.

In order to address these issues, we have decided to implement two hard switches in each office: a light switch and a "smart enable" switch.

The lights can be controlled by a push-button switch on the wall. This switch is wired in a "three-way" configuration, with the other leg connected to a BMS-controlled relay. This configuration provides two important advantages:

- The push-button always toggles the light state, and the response is immediate.
- The lighting control is fail-safe. Even if computer control is disabled due to hardware or software problems, the manual light switch will continue to function normally.

For the lights, we chose to use a push-button rather than a toggle, to avoid the misleading feedback provided by a toggle switch's "up is on" semantics.

The smart enable switch addresses many of the issues raised above:

- If the smart switch is turned off, the office reverts to a known, basic level of functionality. This functionality is provided entirely within the BMS, and no input from the PARC computational network is allowed or necessary. The BMS has been found to be exceptionally reliable at PARC, so we believe that we can use it in this way to guarantee a basic level of functionality.
- The smart switch also supplies power to the occupancy sensor. With the switched turned off, the occupancy sensor is physically disabled. (A red light on the sensor goes out, providing confirmation that it is disconnected.)

Hardware Architecture: Reliability and Flexibility

The model area hardware architecture, shown schematically in Figure 5, was designed to provide an infrastructure that is both reliable and flexible. In order to achieve reliability, we chose to build the basic level of functionality into a commercially available building management system [3]. The BMS, a Control Systems International I/Net system, is based on a networked collection of Z-80 processors.

The processors run a proprietary software system that can be configured to use a wide range of simple control functions. The network allows each processor's outputs to depend on values read or computed by the other processors. The frequencies for sampling inputs and recomputing outputs are also configurable, but the highest frequency possible is once per second.

Access to the network is provided by a *tap*, which provides a simple serial line protocol for reading and writing the processors' inputs, outputs, and intermediate variables, as well as more complex protocols for downloading new control functions, setting other parameters, and uploading collected trend data. Proprietary PC-based software is provided by the BMS manufacturer for end-user network configuration and monitoring.

Our current configuration consists of a single BMS processor, with 96 external inputs and 96 external outputs, together with taps for a PC running the manufacturer's software and a Sun SparcStation-2 running our custom software. (During development, the model area BMS is isolated from the main BMS controlling the rest of PARC, which comprises 19 processors with approximately 800 external inputs and 500 external outputs.)

While the BMS provides a straightforward and reliable way to connect arbitrary input and output devices to our computational infrastructure, each device must be physically connected to a BMS processor. This is appropriate for permanent fixtures such as temperature sensors and overhead lights, but inconvenient and expensive for controlling less permanent appliances, such as computer monitors and desk lamps. For these devices we will use a commercial power control system (X-10) that transmits signals over the AC power lines. This system has two components: a transmitter activated by a computer, and receivers that connect each appliance to the AC wall outlet. Each receiver is assigned a separate address, and by sending appropriate codes across the power line from the computer, individual appliances can be switched on and off.

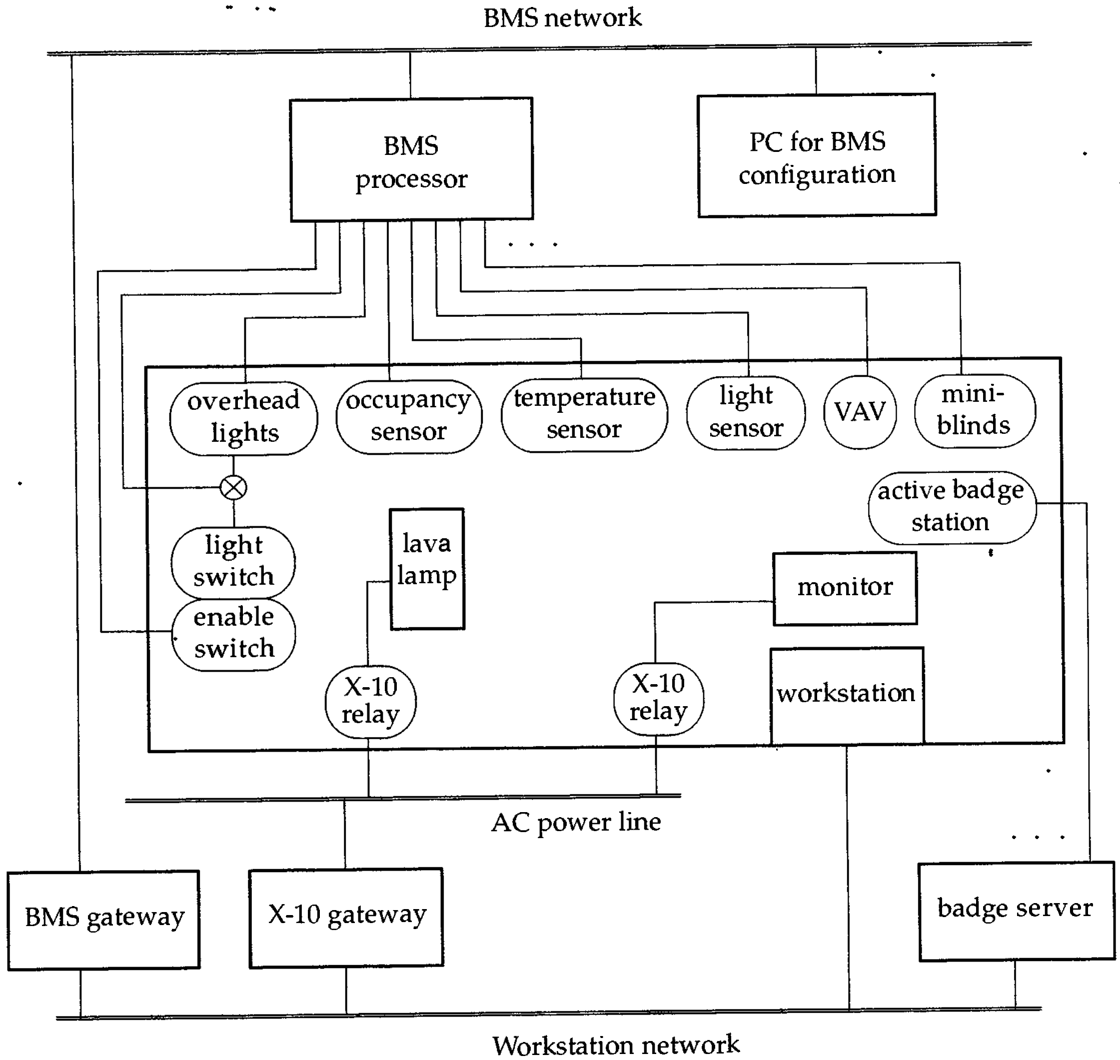

Figure 5: The model area hardware architecture, showing a typical office together with the communications infrastructure

4. Software

The software controlling the model area is shown schematically in Figure 6, and is divided into five areas:

- *BMS internal control functions*: These are implemented by appropriate configuration of the BMS processors, and are responsible primarily for moment-to-moment maintenance of constraints established by higher-level software.
- *High-level control algorithms*: These run on user workstations or servers, and combine information from multiple sources to determine appropriate goals for the low-level control.

- *User interface applications*: These run on user workstations or mobile computing devices, and provide both direct control over the current state of the office and longer-term configuration and preference settings.
- *Gateways*: The gateways provide access to the BMS and power line control systems from PARC's main computer networks.
- *Device control library*: This library provides a simple, uniform protocol for interacting with devices independent of the means used to communicate with them.

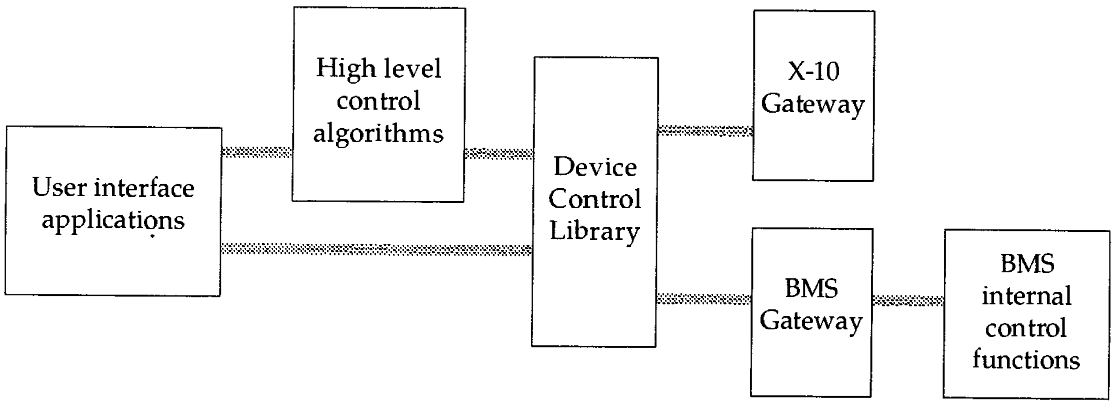

Figure 6: The software architecture

BMS Internal Control Functions

The functions of the internal BMS software are described below:

- Temperature control: Given a setpoint for the desired temperature, the VAV settings are continuously adjusted based on feedback from the room's temperature sensor and other information about the supply air.
- Lighting control: The lighting can be operated in several different modes. A setpoint can be specified for the desired light level, and the BMS internal software will continuously adjust the artificial lights and miniblinds to maintain that light level. Alternatively, the lights and miniblinds may be controlled automatically by higher-level software or by direct input from the user. In addition, the internal BMS software can be configured to automatically switch the lights off when the occupancy sensor indicates a vacancy, and switch them back on when an occupant returns.
- As stated above, the BMS was chosen to be responsible for the basic level of office functionality. When users turn off the smart enable switch, or if a failure of higher level software is detected, the office will revert to a minimum level of functionality that is defined entirely within the BMS. This basic level forces the temperature to a pre-defined standard setting, disables light dimming, and disables automatic switching of lights.

Gateways

There is currently one gateway, which serves as a bridge between PARC's computational network and the BMS. The gateway operates as a network server, exporting a simple RPC interface that allows clients to read and write devices on the network. The primary functions it provides are:

- A simple naming service, mapping a location (e.g. '`ROOM 2130`') and device type (e.g. '`TEMP`') to a BMS device address.
- A simple password-based form of access control, limiting the devices each client can read and write.
- A polling service, allowing a client to specify that an input should be polled at a specified frequency and the client notified when its value goes outside a specified range.
- Arbitration and optimization of requests from multiple clients.

A similar gateway providing access to the X-10 devices is under development.

Applications

We are developing a basic set of software applications that provide control over the office environment from desktop workstations and mobile computers.

Our approach has been to make the smart office feel like the natural extension of a normal office by keeping all of the mechanisms (the BMS and the gateways) hidden wherever possible. From the user's standpoint, an office has a collection of setpoints and other persistent states that are implicitly available to anyone physically present in that office. Access is provided by a small set of X window-based application programs that furnish basic information such as:

- current readings from sensors and actuators located in the office, such as the light level and VAV setting;
- current readings from non-local sensors of potential interest, such as the outside air temperature;
- controls for actuator setpoints, such as the office thermostat; and
- controls for occupant preferences, such as manual vs. automatic lighting.

We are also developing applications that graphically display energy costs. We believe that this will increase energy awareness and lead to more informed energy consumers.

We plan to use mobile ubiquitous computers to control office environments in novel ways. One software application will enable hand-held PARCTab computers [1] to be used as remote control devices for office lights and temperature. The location of a mobile Tab can be inferred from the location of the stationary base station with which it communicates. This will allow the Tab application program to identify the appropriate sensors and actuators, and to adapt its display appropriately as the Tab is carried from room to room.

If based on occupancy data alone, setbacks have serious limitations. As an example, some devices (such as light bulbs and disk drives) may fail prematurely if they are power cycled too frequently. Therefore, it is not necessarily cost effective to automatically turn devices off the instant that a room becomes vacant. In general, the optimal strategy will depend on the particular appliance in question and on how long the office is going to be vacant. In the case of air conditioning, long response latencies make it impractical to set back individual offices without some information about the occupant's probable time of return.

For these reasons, we are investigating the use of various sources of information to predict when rooms will become vacant, and when the occupant is likely to return. Some of the potential sources of such information are:

- user-supplied occupancy profiles;
- personal calendar managers, showing both when the occupant plans to be at work as well as when meetings are scheduled;
- statistical analysis of occupancy patterns; and
- active badges or other location-sensing technology (for example, going to a conference room or leaving the building may offer clues as to the probable length of absence from an office).

5. Privacy Issues

Serious privacy issues result from the placement of remotely-readable occupancy sensors in individual offices and from the use of active badges. These issues are even more serious if the occupancy data is trended or analyzed for use in the manner suggested in the preceding section. In our model area, occupancy data will in principle be available to any desktop computer via the regular local area network. Office occupants have very legitimate concerns about how such information might be used.

Ideally, we would like to be able to use occupancy data for purposes like energy conservation, while at the same time protecting individual privacy. If occupancy sensors or active badges are to be used, we believe that privacy can only be protected by a combination of technical and social means.

As one step toward developing a social consensus that limits the misuse of such data, we have drafted a model agreement in the form of a contract between office workers and senior management. The agreement lays down the terms under which such occupancy data can be used. In our case, the office worker agrees to the installation of an occupancy sensor or use of an active badge provided that the information obtained from them is used only for the limited purposes of building energy and comfort management.

As described above, we have also attempted to address the privacy issues in the design of the office hardware. The smart enable switch gives the user the ability to physically disable the occupancy sensor at any time. This gives the office worker a simple non-confrontational way to disconnect the sensor if they feel their privacy concerns are not being adequately addressed.

Due to the centralization of occupancy information, the current hardware architecture cannot provide absolute assurances of privacy. In the long term we would like to redesign the hardware so that occupancy data is available only for local computation in the office and is not transmitted over the network.

6. Summary and Future Directions

The energy audit that we conducted suggests that substantial savings can be achieved through independent control of individual offices. Much of the hardware needed to save this energy can also be used to give occupants a greater level of personal control. This substantial overlap suggests that making office environments more responsive to individual preferences and activities can serve both to enhance comfort and save energy.

As the hardware and software installation nears completion, we will soon be in a position to evaluate this hypothesis and quantify the costs and savings involved. Our plans include: collection of detailed sensor data from our model area offices, evaluation of the effectiveness of various energy saving strategies, and experimentation with novel control strategies.

Finally, we plan to quantify the cost of sensors, wiring and other hardware, and to demonstrate how occupant comfort can be increased and energy saved in a cost effective manner.

7. References

1. Weiser, M. *The Computer for the 21st Century.* Scientific American 265(3), Sept. 1991, pp. 94–104
2. Loveday, D. L. and Virk, G. S. *Artificial Intelligence for Buildings.* Applied Energy 41, 1992, pp. 201-221
3. Control Systems International, Inc., Dallas, Texas
4. Energy User News, 17(5), May, 1992
5. Want, R. and Hopper, A. *Active Badges and Personal Interactive Computing Objects.* IEEE Transactions on Consumer Electronics 38(1), February 1992, pp. 10-20