

Well-connected microzones for increased building efficiency and occupant comfort

*Michael P Andersen, Gabe Fierro, Sam Kumar, Joyce Kim, Edward A. Arens, Hui Zhang, Paul Raftery, and David E. Culler
University of California, Berkeley*

ABSTRACT

Thermal microzones, such as those created by Personal Comfort Systems (PCS's), have been shown to be capable of maintaining occupant comfort in buildings despite large zone temperature deviations from recommended "comfortable" temperatures. We show that by developing well-connected microzone devices that report real-time telemetry and respond to programmatic actuation requests, these established capabilities can be synergistically combined with occupant-aware building applications to enable new methods of comfort and energy efficiency maximization. We present a novel digital controller for a well-studied (previously analog) PCS chair that maintains 90% occupant comfort over a 20°F range while using less than 16W for heating and 3.6W for cooling. Using the digital PCS chair, a furniture microzone system is evaluated in two settings. The first utilizes a controlled testbed to demonstrate the functionality of the microzone control system along with a representative set of control algorithms including comfort-driven setpoint manipulation and transparent demand response. The second is a two month pilot study of 30 digital chairs given to occupants of a building that exhibited low occupant comfort levels. The chair telemetry was used to analyze occupant comfort in real-time, and drive HVAC control. We conclude with a brief overview of the questions posed by this platform, along with research areas enabled by such fine-grained real-time control and measurement of building occupant comfort.

1. Introduction

In building engineering, occupant comfort is addressed by defining a range of temperature and humidity that are considered to be comfortable and by maintaining the indoor environment conditions within this range. It is assumed that stricter adherence to the comfort zone will yield higher occupant comfort. The problem is that occupants are not comfortable; despite massive energy expenditure – 10-20% of grid energy in developed countries (P'erez-Lombard et al. 2008, 394-398) – a large portion of occupants remain dissatisfied with the buildings' thermal environments. A survey of 215 buildings in North America and Finland (Huizenga et al. 2006) showed that 42% of occupants are dissatisfied with the temperature of their environment, compared to only 39% who are satisfied. In contrast, ASHRAE standards set the target thermal satisfaction rate at 80%, a goal which only 11% of surveyed buildings manage to achieve (Zhang et al. 2015; Huizenga et al. 2006). This paper introduces a rich new paradigm of cooperative control between personalized micro environments and building control systems that opens up several new approaches to addressing both occupant discomfort and building power consumption. While personalized micro environments have long been possible using

traditional appliances such as space heaters, in general such environments conflict with the building environmental control and are easily overpowered. Indeed, the thermal load from a space heater is indistinguishable from the unwanted heat of other electrical appliances that the HVAC system is engineered to remove. The resulting ineffectiveness and inefficiency of these micro environments has thus far limited their usefulness. This situation has changed with the emergence of a family of Personal Comfort Systems (PCS's), that allow for significant modulation of the perception of thermal comfort while inducing little actual heat transfer. Such systems do not significantly couple into the overall control system and therefore do not “fight” with the building. Leveraging this technology and combining it with mobile devices and wireless sensor networks enables a new paradigm of connected microzones. This not only enables a variety of occupant-centric control methodologies that allow greater optimization of occupant comfort and building energy consumption, but it opens up avenues for cooperative control that places the human “in the loop”. Sections 2 and 3 introduce the design and requirements of a connected microzone system, and the building applications that leverage it. Section 4 introduces PCS's – an enabling technology for ideal microzone construction. Using this technology, we design and implement a well-connected furniture microzone that forms the basis of a cooperative microzone architecture. Using a controlled chamber study, Section 6 validates the technical feasibility of the connected microzone system, demonstrating cooperative control between the building management system and the microzones.

2. Connected Spaces

Within buildings, a large part of occupants' satisfaction comes from their thermal comfort, affected by their immediate environment. Sometimes this environment can be operated in partial isolation from the rest of the building, maintained by devices such as space heaters, fans, small window air conditioners and – as discussed later – PCS's. The choice of device greatly affects how isolated these microzones are, with space heaters and air conditioners forming only a crude approximation to a microzone, having a large impact on the macrozone maintained by the building HVAC system, as shown in Figure 1. While not all buildings have these microzones, it is well established that having this finer granularity in environmental control improves occupant comfort, especially if the microzone is controllable by the occupant (De Dear et al. 2013). First, we consider how connected microzones can be leveraged to improve occupant comfort and building efficiency even further, with the design choices involved in microzone implementations deferred to Section 4 and 5.




Figure 1: A person interacting with a microzone inside the building controlled macrozone

2.1 Increased visibility

Despite the comfort benefits, microzones are not common. One of the principal arguments against occupant-controllable microzones, especially crudely formed ones that influence the macrozone such as those defined by space heaters, is that the improvement in occupant comfort does not justify the very probable inefficiencies that result from their opacity (Arens and Brown 2012). If the HVAC system is not aware of the microzone, the macrozone and the microzone could conflict with each other. Consider the case where a cold occupant turns on a space heater. The room thermostat then detects a rise in temperature, causing the building to dump more cold air into the room. The occupant responds by turning up their space heater and the process continues, wasting energy. Even if the microzone did not impact the macrozone as heavily, continuous simultaneous heating in the microzone and cooling in the macrozone is not ideal. These situations occur due to a lack of visibility: the building does not know about the heating microzone, nor the cold occupant. It acts as best it can, given the limited information available to it. The solution to these problems, is to connect the microzone to the building. If the devices that maintain microzones are part of a network that is integrated with the building control applications, the systems can not only be prevented from fighting each other, but they can synergistically combine to become both more efficient maintaining occupant comfort and minimizing energy use. Well-connected microzone devices must expose their current configuration over the network. This allows the processes controlling the macrozone to factor the microzones into their calculations. Furthermore, the microzone must be remotely controllable. This allows the building to push back on a microzone when it knows of a more efficient solution. Consider a scenario where all the microzones in a zone are too cold and have active space heating: it is more efficient for the building to change the temperature of the whole zone and then override the microzones to turn off the space heaters. Finally, once microzone devices are connected to the building, there is a very small incremental cost to including additional sensors. Temperature and humidity sensors can be added to provide the exact environmental conditions at a location near each occupant. Occupancy sensors can be added to allow the building to vary control logic depending on how many people are in an area, for example reducing airflow while keeping within minimum airflow regulations.

2.2 The interaction model

When microzones become integrated into the control logic of a building, a design choice regarding the interaction model needs to be made. At present, there are roughly two categories of methods for occupants to control their thermal environment. The first is comfort-centric: occupants report their comfort level, and some process acts on this to improve their comfort, for example an email to the building manager that results in a setpoint tweak. This is also the principle of systems such as Comfy (Krioukov and Culler 2012) and ThermoVote (Erickson and Cerpa 2012) that collect and act on comfort data autonomously. The second category is setting-centric – occupants have the ability to change a setpoint directly, or to enact changes in HVAC settings directly such as calling for a brief blast of cold air. This is not very scalable, as typically these settings impact more than one person and so require an arbitration layer in front to remain fair. By introducing this arbitration layer, the system begins to fall more into the comfort-centric category. We introduce a third category of interaction: microzone-centric. Here, users interact with the devices in their microzone directly and have complete control over them. Higher level comfort information is inferred from these settings.

The advantage of a comfort-centric model is that the building may know of the most efficient solution to the problem. For a cold occupant it may be that decreasing the airflow into a room is a better solution than turning on a space heater, but the occupant may not be aware of that option. The disadvantage is that sometimes the resulting action is not what the user wanted or expected. The setting-centric model gives the user control, but is not feasible at scale as it cannot resolve differing comfort requirements in the same space. In a microzone-centric model occupants interact directly with their microzone, which reacts immediately, giving them the psychological benefits of feeling in control, as in the setting centric model. Thermal comfort is a “state of mind” and the psychological aspects of feeling in control contribute significantly to this (De Dear et al. 2013). By inferring the occupants’ comfort from their interactions and using this in the control logic, a hybrid model can be formed that changes strategy from using the microzone to using whatever is most efficient. This transition can happen once the user becomes comfortable or the macrozone has finished adapting, for example. Furthermore, to participate in a microzone-centric model, occupants do not need to change their behavior from how they interacted with unconnected microzone devices. If they wish to turn on the desk fan, they can simply physically turn on the fan.

2.3 System Architecture

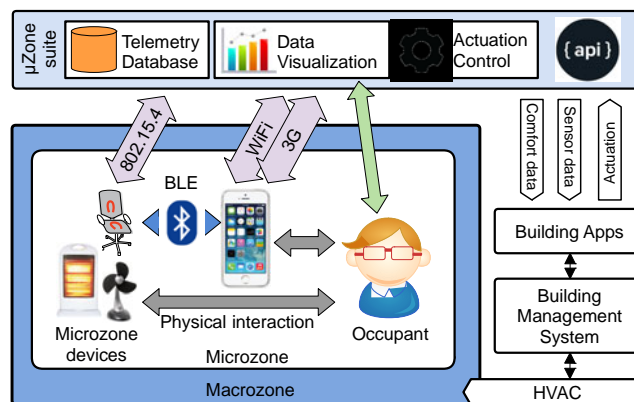


Figure 2: A microzone-aware building control system

Figure 2 shows the architecture of our system using the hybrid interaction model; the occupants interact with the microzone control devices directly. For some devices this is a direct physical interaction, whereas more advanced devices such as the PCS chair, described later, are controlled via a mobile application that connects to the device via Bluetooth. The telemetry flows to a suite of microzone-interfacing software where it is archived and made available via a data visualization service to building occupants and operators. In addition, building applications – such as those implementing control algorithms – can use the microzone control suite to subscribe to real-time sensor and comfort streams, or query historic data.

3. Building Applications

The introduction of connected microzones drastically changes the model of a building as seen by control applications. With devices streaming information from distributed temperature

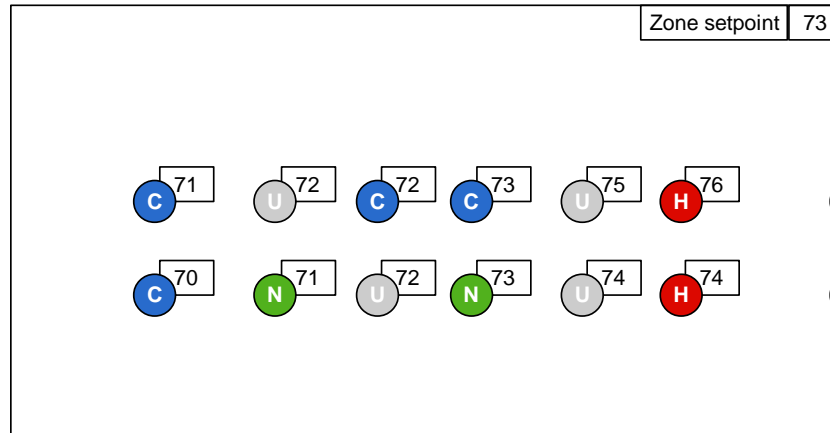


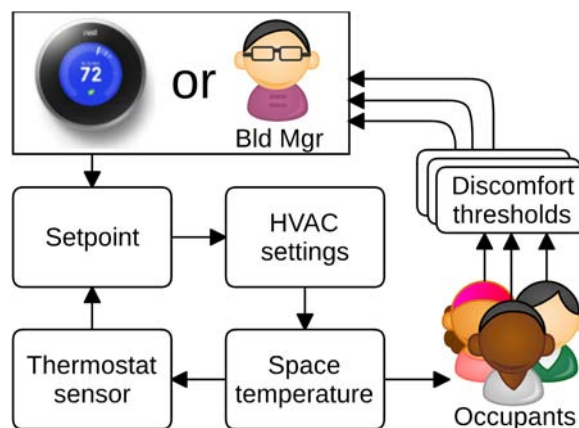
Figure 3: A shared space containing connected microzones displaying comfort, occupancy and temperature information as seen by the building control software. Letters denote **C**old, **N**eutral and **H**ot occupants.

and occupancy sensors, as well as the inferred comfort information from microzone device usage, the view of the building resembles Figure 3. Here, a building application can see where people are, how they are feeling and the temperature of local environment. This capability, combined with the finer granularity in control, allows for a new class of building applications. These applications can assist with HVAC control – dynamically changing setpoints and airflow based on comfort, temperature and occupancy – as well as enable “smart building” behavior that improves occupant experience in other areas.

3.1 Advanced HVAC control

Consider Figure 4, a representation of the thermal environment feedback loop present in most buildings today. Occupants work in a shared indoor climate with no personally controllable buffer. If they become sufficiently unhappy with that environment, they may change the setpoint (if possible) or complain to the building manager. Each person has a different ideal environment and a different threshold at which they are willing to express their discomfort. The setpoint is then used in a control loop whose input is the temperature reading from the zone thermostat.

This system has several well known problems. The building manager (if present) has to



determine the optimum setpoint given only a subset of comfort information that has been biased by how willing an occupant is to complain. This leads to a less than optimal choice. If there is no building manager then the setpoint will only reflect that last person's comfort requirements. As a large zone with multiple occupants may only have a single temperature sensor, the HVAC system is only optimizing for a single point in the zone, and not necessarily one that represents the midpoint of temperatures across the zone (Arens and Brown 2012).

These problems can be solved by introducing connected microzones. Figure 5 shows an alternative scenario where occupants have a buffer, maintained by microzone devices that can alter their thermal perception of the environment. Every occupant can match their microzone to their own comfort requirements and in so doing generate telemetry that is used by an arbitrating




Figure 5: A new comfort feedback loop. Occupants remain comfortable in microzones and telemetry is used to automatically adjust HVAC settings

building application to maximize comfort autonomously and objectively. The HVAC feedback loops are based on distributed temperature data, which greatly increased the accuracy over a single acquisition point. Even if a building does not take advantage of the many other applications enabled by connected microzones, simply introducing a microzone buffer and using its telemetry to repair the feedback loop would drastically increase occupant comfort (Arens and Brown 2012).

3.2 Enabling the empathetic building

In addition to improved HVAC processes, the information produced by connected microzones enables a building to learn about its occupants. Many microzones are occupied by the same person every day – an office for example – or allow for identification of the occupant via the method of interaction – a mobile phone, for example. If the building application studies the microzone interaction telemetry and correlates it with other streams, such as the sensors in the microzone and outside weather, it can begin to determine the relationship between measurable variables and an occupant's comfort, forming a profile. These profiles can be used to predict the comfort of a person, given information about the environment or, conversely, to predict the ideal environment to ensure a given person's comfort.

Consider a person who has just cycled in to work and sits down at his/her desk. The building knows that when the weather is decent (and they can cycle), this occupant always comes in to work feeling hot and so it can turn on a fan automatically. Taken further, the building can

learn that this person always arrives at around 8AM, so it can pre-cool the office by a few degrees before the occupant even arrives. Multiple profiles can be used together, for example the building can read the shared calendar and know that at 1 PM a specific meeting room will be occupied by a given set of people. It can heat or cool the room in advance so as to maximize the average comfort of the occupants according to their profiles.

While many of these ideas have existed in the past, the lack of available information has made them difficult or impossible. With well-connected microzones, these become simple software applications.

4. Thermal Comfort

Crude approximations to microzones can be implemented using traditional devices such as space heaters and windowmount air conditioners, but these devices are not energy efficient and affect neighboring occupants. For the most part, they are working on the same principle as larger HVAC systems: heat up or cool down large volumes of air in order to maintain occupant comfort. In order to implement microzones efficiently, enabling long term deployment that yields energy savings in addition to the increase in occupant comfort, it is necessary to use microzone devices that manipulate comfort in a different manner.

4.1 The shifting goalpost problem

The disparity between the energy spent on controlling ambient temperature and occupants' dissatisfaction with their environment can be partially attributed to the distinction between air quality – a function of temperature, humidity and freshness – and thermal comfort. Buildings attempt primarily to achieve occupant comfort by keeping the indoor air fixed at a constant temperature and humidity (Brager et al. 2015). Occupant comfort, however, is a moving target. Comfort requirements vary among individuals according to factors such as age, gender, body mass, metabolism and thermal adaptation (Zhang et al. 2015). The thermal preferences of even a single individual may vary according to outdoor climate and clothing factors (De Dear et al. 2013). The diversity in occupants' thermal comfort requirements limits the number of satisfied occupants to about 80% when the ambient temperature of all occupants is maintained at a fixed setpoint (Arens et al. 2010; Zhang et al. 2011).

In addition to the differences in people's ideal temperatures, psychological aspects of the environment play an important role in determining comfort level. As an example, occupants of naturally ventilated buildings are comfortable over a wider temperature range due to perceived control over their environment (De Dear et al. 2013). Air flow also directly impacts thermal comfort and has been shown to widen the range of temperatures in which an occupant feels comfortable (Zhang et al. 2015; Zhang et al. 2011). This suggests that it is possible to adjust people's perceptions of their ambient environment, and thereby affect their thermal comfort levels, without actually changing their temperature.

4.2 Personal Comfort Systems

One alternative approach to achieving higher levels of occupant comfort works by manipulating the occupant's perception of their environment without significant heat transfer.

This idea is well-studied in existing building architecture and HVAC literature, and is known as Task-Ambient Conditioning or, more recently, as Personal Environmental Control (Paciuk 1989; Glicksman and Taub 1997; Falkner et al. 1999; Amai et al. 2007; Bauman and Arens 1996; Zhang et al. 2015). These systems take advantage of the fact that the range of temperatures within which a person will remain healthy is substantially wider than the range of temperatures within which they are naturally comfortable, so there are no adverse effects when manipulating the human body into feeling comfortable outside of the standard comfort zone but within the healthy zone. Studies (Zhang et al. 2015) have shown that PCS's can manipulate perceived ambient temperature by several degrees Celsius in either direction, enough to significantly extend the perceived comfort zone.

As they operate at the occupant level, Personal Comfort Systems allow each individual to meet their thermal comfort needs without significantly affecting the thermal environment of other nearby occupants. This leads to higher occupant comfort by allowing heterogeneous comfort requirements to be simultaneously met.

4.3 Reducing building energy

Table 1: Comparison of maximum per-person operating power requirements for different devices influencing thermal, and Perceived Temperature Difference (PTD) produced by PCS [18][14][16][19] *Includes effect of 30.8 W footwarmer

Device	Cooling	Heating	PTD
Central HVAC	500W	500W	
Space Heater		400W	
Chair and Desk Fan	4.8W	16W	-4K / 7K
Footwarmer		21W	2.2K
Heated Keyboard / Mouse		28.6W	6.5K*
Head/Hand Ventilator	2W		-5K

Microzones in general can keep occupants comfortable over a wider setpoint deadband, but this is not guaranteed to reduce building energy consumption unless the microzone devices are more energy efficient than the building systems. With microzones based on Personal Comfort Systems, significant energy savings can now be realized. A 2014 study detailing parametric simulation of six distinct building types in seven different ASHRAE climate zones (Hoyt et al. 2015) showed that widening the setpoint deadband decreases the total energy spent on HVAC by roughly 10% per degree Celsius in either direction. As Personal Comfort Systems can adjust temperature perception by several degrees (the PTD column in Table 1), the HVAC energy savings obtained by relaxing the regulation of ambient temperature come at no cost to occupant comfort (Zhang et al. 2015) and greatly outweigh the cost of operating a system per occupant (Hoyt et al. 2015).

While these systems offer tremendous potential for improving occupant comfort while simultaneously decreasing energy usage in buildings, so far they have operated independently from the building's environmental control. When these Personal Comfort Systems become

digitally controlled and are connected to the building’s HVAC system, they become the ideal connected microzone capable of influencing a person’s thermal comfort while using minimal energy and not influencing other occupants.

5. Design and Implementation of a Furniture Microzone

To realize a PCS based microzone with the requisite connectivity and remote actuation capabilities, it was necessary to choose a PCS and design an advanced controller for it. A PCS chair was chosen as it is easily deployable, offers very reliable occupancy detection, affects only a single occupant, uses a maximum of 16W and has been shown in human trials to be extremely capable in manipulating an occupant’s perceived comfort, enabling over 90% of occupants to be comfortable over a wide range of 11 degrees Celsius (Pasut et al. 2015). This section details the design and implementation of a PCS chair based furniture microzone.

The next-generation PCS chair is composed of three tiers, as shown in Figure 6, and interfaces with an Android application running on the user’s phone. The combination of phone and PCS chair fits into the connected microzone and building application architecture described in Sections 2 and 3.

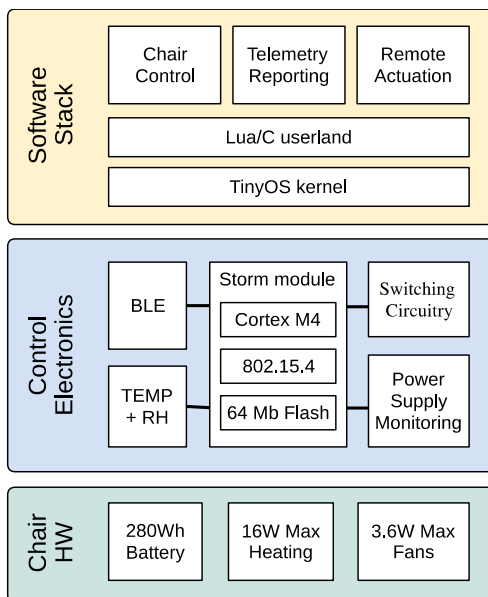


Figure 6: The Personal Comfort System chair platform




Figure 7: The PCS chair hardware showing the heating and cooling capabilities

5.1 Mechanical design

At the base, we are using the same physical chairs that have been used in previous studies (Pasut et al. 2015; Zhang et al. 2015). The mechanical design of the chair is shown in Figure 7. It consists of two heating strips and three fans, installed in a mesh-type office chair. The heating strips, one on the bottom and one on the back, are in contact with the chair occupant. The fans, two on the bottom and one on the back, are situated further back from the mesh so that they increase airflow over the surface of the occupant when active. The heating strips use a maximum

of 16 W, while the fans use a maximum of 3.6 W allowing the 280Wh battery to power the chair for several days of average use. Placed in the bottom of the chair is a contact switch which closes when the user sits down, providing accurate occupancy information. As an energy optimization technique, this switch is also used to power down parts of the chair when it is not occupied.

The 16 W delivered by the heating strip is not powerful enough to significantly affect the occupant's core body temperature, but it is delivered over a fairly narrow area. This causes localized heating which causes an increase in perceived temperature far in excess of what 16 W could cause if distributed evenly over the surface of the occupant. It is this effect that allows the chair to be so energy efficient. Likewise, the fans do not cool the air down but rather increase the airflow over the surface of the occupant, which increases natural convective and evaporative cooling. Human trials (Pasut et al. 2015) have indicated that this chair is capable of keeping 92% of subjects comfortable over a range from 18°C to 29°C (64.4°F to 84.2°F).

5.2 Control circuitry

Previous designs (such as those used in the aforementioned human trials) utilized an analog controller mounted to the side of the chair with a switch for heating/cooling and a rheostat to control intensity. This is replaced by a digital controller – the middle layer in Figure 6. At the core of this controller is a Storm – a reusable Cortex-M4 module running a TinyOS based software stack. The Storm module provides an IEEE 802.15.4 radio so that chairs can send telemetry to the control suite described in Section 2, as well as have their settings changed remotely. The carrier board provides application specific functionality – a monitoring power supply, interface circuitry and Bluetooth connectivity.

The Storm module is used due to its ultra-low operating current. Although the power requirements of the fans and heating strips dominate, they are used only periodically. The rest of the controller is continuously on and sending telemetry so it is advantageous to minimize its quiescent current. By using the Storm, a chair that is unoccupied or occupied by a person comfortable with his/her environment consumes less than 1mA on average while sending telemetry packets once per second. With the equipped battery, this allows the chair to continue sending telemetry for years between charges. Although the use of a more powerful core module such as a Raspberry Pi or other Linux machine is possible given the large battery pack, these devices consume hundreds of mA at a minimum and would reduce the maximum time between charges by two to three orders of magnitude. A power monitoring circuit monitors the battery to alert the operator when the chair needs recharging.

The interface circuitry connects to the heating strips and fans to provide fine-grained setting of the fan and heat intensity independently on the seat and back of the chair. The heating strips use an energy efficient pulse width modulated (PWM) open drain circuit so that all the energy is dissipated in the heating strip, irrespective of intensity setting. This is an improvement over the analog rheostat control, which wasted energy in the control circuitry. Unfortunately, as we were retrofitting existing chairs, the fans were not compatible with PWM control – they did not have a dedicated PWM control line and attempting to PWM the power line resulted in audible buzzing and motor stalls. We therefore use voltage mode control to modulate the fan intensity. As the fans use only 3.6W (and substantially less as the voltage drops off) this is acceptable, although future revisions of the chair will use PWM-compatible fans such as those from personal computers.

A temperature and relative humidity sensor was added to the controller to enable accurate distributed environmental monitoring. This information enables direct comparison between the settings that the user chooses, and the environment that he/she is in. In actual deployments, the temperature at the location of individual chairs in a room can differ significantly from the temperature reported by the thermostat. The temperature sensors in the chairs can be used to compensate for this in HVAC control loops. The sensor is thermally isolated to minimize heat gain from the rest of the controller, a design choice we found lacking in the majority of smart thermostats available off the shelf.

5.3 Reliable telemetry delivery

While real-time telemetry is important for control applications, there are also applications that rely on historical data. The building application that forms comfort profiles about its occupants is an example of such an application. This means that even if we are unable to report real-time telemetry (if the chair loses Internet connection or the microzone services are down), it is worth persisting the telemetry on the chair and transmitting it when connectivity is restored. To do this, the software stack (Tier 3 in Figure 6) writes telemetry to a log on the Storm’s flash chip.

In addition to log replay over 802.15.4, there is an additional improvement to redundancy when there is a mobile phone attached via Bluetooth: the chair will replay the log via the phone’s Internet connection. This capability enables the chairs to be deployed entirely without 802.15.4 infrastructure, although there will be delays in data delivery when the users remove their phones from the chair’s vicinity. This is the configuration we have used for small studies so far, as it significantly reduces the deployment effort and allows chairs to be moved around without worrying about reaching an 802.15.4 router.

6. Experimental Verification

In order to verify and demonstrate the functionality and technological feasibility of our connected microzone system, a representative subset of building applications are used to conduct a series of full-coverage tests. The tests are performed in a self-contained chamber (Figure 8) that can be manipulated programmatically in a controlled and repeatable way. The chamber contains two PCS-chair-based microzones within a macrozone that is maintained by a controllable space heater and window-mount air conditioner. To emulate typical BMS control of the macrozone, a




Figure 8: The isolated test chamber used for technology verification and control application testing

building application listens to a thermostat within the chamber and actuates the air conditioner when it calls for cooling. The space heater is used to simulate heat gain. The advantage of this isolated chamber over placing the microzones in a standard room and using existing HVAC control is that we can rapidly create precise edge conditions that would otherwise be difficult or slow to emulate.

6.1 Setpoint setback

We validate the ability of the system to allow a building to respond to microzone interactions with a simple building application that switches between two setpoints based on the average fan settings of the PCS chairs. The logic is as follows: if there are no occupants in a zone, or the occupants are comfortable, then a relaxed setpoint is used. As soon as occupants feel uncomfortable, the setpoint changes to a more conservative setting. In an actual deployment, this type of application can be used to take advantage of the fact that a person who is coming into the building from a warm climate may have already compensated for the temperature by modifying their clothing, and may not require the environment to be as cold as traditional building standards require. If they do become uncomfortable, however, the building does respond.

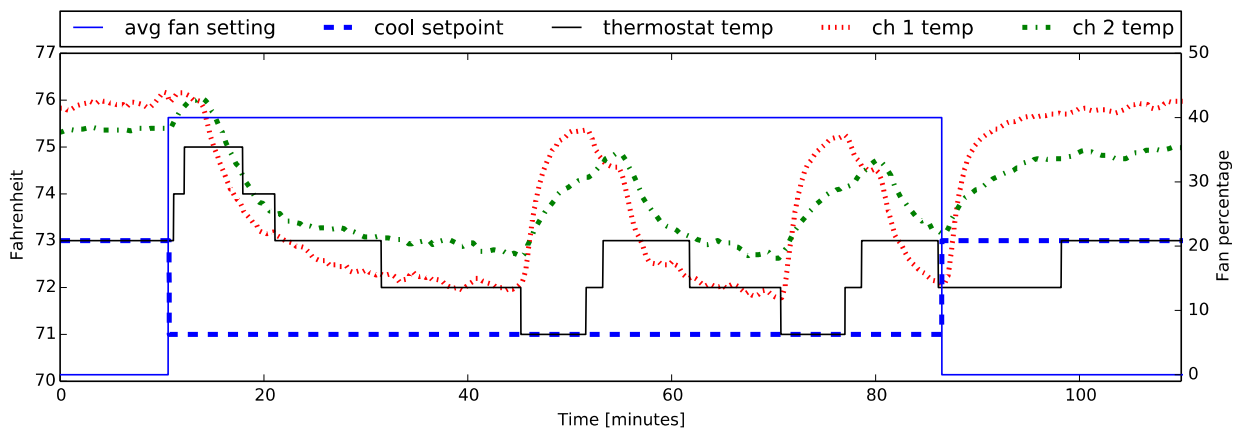


Figure 9: A simple responsive environment example where the setpoint is altered based on the average occupied chair fan setting

Figure 9 shows the telemetry obtained from this experiment. We see that the average fan settings of occupied microzones (the solid blue line) remain at zero until 10 minutes when it spikes to 40%. It remains there until 90 minutes in, where it drops down to zero again. From the point of view of the application, an unoccupied chair is the same as a chair occupied by a comfortable person. The dashed blue line shows the macrozone setpoint. We can see that the building is responding to the microzone interaction by lowering the setpoint by two degrees during the interval that the fan settings are nonzero. The narrow black line indicates the temperature reported by the thermostat. When this is higher than the setpoint, the air conditioner in the chamber turns on. The rapid increase in temperature after the air conditioner turns on can be attributed to sudden mixing of warm air that has risen within the chamber.

While simple, this example shows how a building might respond to indirect information from the occupants' interaction with their connected microzones. As the users adjust their fan settings, they immediately feel cooler, and the building slowly adjusts the environment around them to increase their comfort.

6.2 Bi-directional interaction

The previous experiment validated that a building application can be made to dynamically reconfigure the zone according to information from the connected microzones. It is necessary to further validate that the system is capable of supporting the reverse direction: the building controlling the microzone.

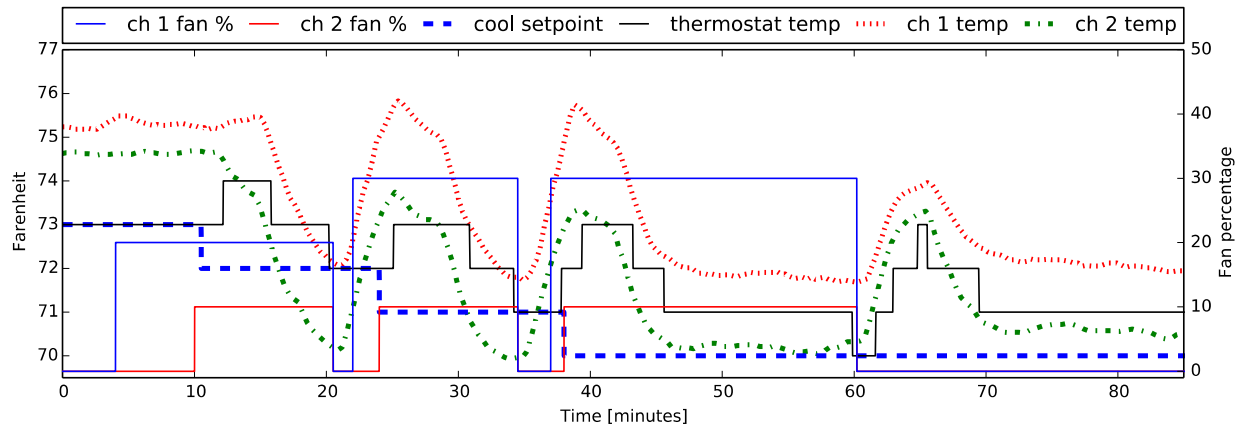


Figure 10: A bi-direction interaction between microzones and the building: the setpoint is adjusted whenever a majority of occupants are uncomfortable, and the chairs are automatically reset when the new temperature is reached

To do so, a more complex building application was used. Here the logic waits for a majority of occupied chairs to indicate that they are too cool or too warm. When that occurs, it adjusts the setpoint down (or up) by a degree and waits for the room thermostat to register the target temperature. When the target temperature is reached, the building actuates the chairs, sets their settings back to neutral and waits for a new majority. This building application is an example of the hybrid interaction model discussed in Section 2.2.

The telemetry from this experiment is shown in Figure 10. The chamber configuration is the same as the previous experiment: it is emulating a room with strong solar gain. Proceeding chronologically, we see that the first occupant (solid blue) adjusts the microzone fan to 20% at 5 minutes. At 10 minutes in the second occupant also turns his/her fan on, to 10%. At this point, the majority of users in the macrozone are uncomfortable in the same way, so the zone setpoint (dashed blue) is reduced by one degree. The user's comfort is maintained by the microzone devices as the macrozone slowly adjusts. At 20 minutes, the thermostat reading (solid black) reaches the setpoint. Now, the system uses the bi-directional capabilities of the system to reset the microzone devices. The occupants are not yet satisfied, so at 22 and 24 minutes we see the fan settings increase again, and the process repeats. At 60 minutes, after the setpoint has been adjusted down a total of three degrees, the occupants are comfortable, and occupant comfort has transitioned from relying on the microzone devices to relying on the macrozone.

In both Figure 9 and Figure 10, the temperature at the location of the chairs is significantly different from that registered by the thermostat, and each other. It is clear that – as discussed in Section 3.1 – the HVAC control loop would benefit from using the distributed temperature data from the chairs rather than the temperature at the thermostat. Doing so uses functionality already tested and presented here, so is omitted for brevity.

These two experiments are an effective test of the entire system (depicted in Figure 2). The PCS chairs are streaming data to the telemetry archiver in the cloud. The building

applications that encapsulate the thermostat to air conditioner binding and fan-setting to setpoint binding are consuming data from the microzone services and actuating devices in the microzone accordingly. Finally, the figure was generated using a data query from the visualization service.

6.3 Preemptive action

The previous two experiments have validated scenarios where the stimulus for the control logic is the user's comfort level as inferred from their interaction with the microzone. In addition to this, a well connected microzone can also react to stimulus from the building. This reaction can be transient, such as compensating for demand response events, or long term, such as compensating for deliberate changes of the setpoint deadband over the course of a day. The building can alter the environment significantly while using the microzones to transparently maintain occupant comfort. In the case of the PCS-chair microzone, the energy costs of maintaining occupant comfort are far smaller than if the whole zone were kept at a comfortable level (Pasut et al. 2015; Zhang et al. 2015), so maintaining occupant comfort does not undermine the goal of energy savings.

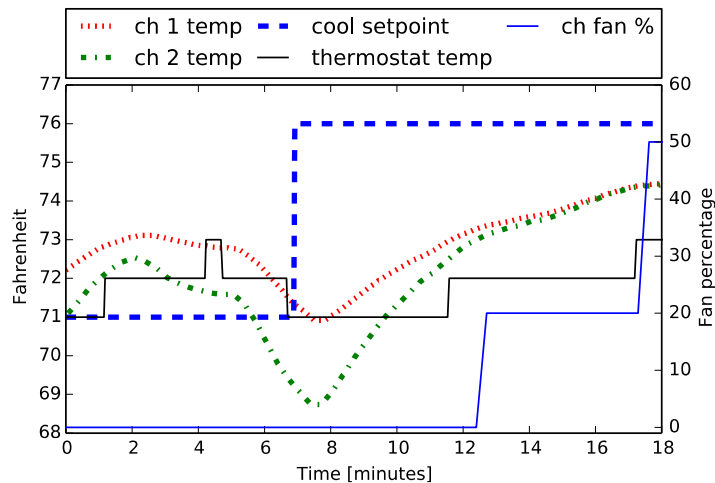


Figure 11: Automatic PCS chair fan actuation in response to a DR event

Figure 11 shows the telemetry from a building application implementing transparent demand response compensation. The building receives a demand response signal at 7 minutes, and increases the zone cooling setpoint to reduce power consumption. It begins monitoring the PCS microzones, and increases the fan settings at 13 minutes and 17 minutes as the temperature rises. This experiment demonstrates the system performing automated compensation for deliberate changes in the environment.

The same technique – monitoring the local environment and automatically actuating the microzone to compensate – can be used over longer periods of time to reduce the perceptibility of widened setpoint deadbands. Prior work (Hoyt et al. 2015) has established that considerable energy savings (10% per degree Celsius) can be obtained by relaxing the setpoints, and PCS chairs have been demonstrated to allow this without decrease of occupant comfort (Pasut et al. 2015), but so far this combination has only been tested over short periods of time. The system has required the user to continuously adjust their chair settings as the environment floats within the widened deadband. While we do not yet have the data to determine if this requirement affects

occupant satisfaction in long-term deployments, this experiment demonstrates how the system can use automated actuation to reduce the required user interaction. Manual changes need only be made when occupant comfort requirements deviate from their normal patterns.

7. Conclusion

This paper introduces the concept of connected microzones that work in cooperation with the building control system. We show how such a system would be architected, and how building applications can use connected microzones to improve the efficiency of HVAC control. We discuss the nature of thermal comfort, and the energy consequences of adjusting the temperature deadband, showing how PCS's can maintain user comfort over a wide range using minimal power and performing negligible heat transfer. We utilize this characteristic of Personal Comfort Systems to construct an efficient furniture microzone that completes an instantiation of a connected microzone system allowing for viability testing. We perform a small set of controlled experiments, chosen to maximize coverage of system facets and to demonstrate key concepts, confirm its feasibility and potential.

There are several PCS devices other than the PCS chair, which could form effective microzones. Section 4 tabulates a subset of these. Multiple of these devices can be configured to operate cooperatively within a single microzone, which could yield improved results. We have explored two means of interacting with the microzone – a touchscreen on the device, and a mobile application – but there are many other interaction techniques such as gesture control, voice control or inferred control that could be applied to improve the seamless integration of microzones.

It is clear that well connected microzones offer many potential advantages and are worth exploring further. We believe that this platform is a critical step in the evolution of both personal comfort systems and microzones and are confident that it will prove valuable as the field develops.

References

- H. Amai, S. Tanabe, T. Akimoto, and T. Genma. Thermal sensation and comfort with different task conditioning systems. *Building and Environment*, 42(12):3955–3964, 2007.
- E. Arens and K. Brown. Broken information feedback loops prevent good building energy performance—integrated technological and sociological fixes are needed. 2012.
- E. Arens, M. A. Humphreys, R. De Dear, and H. Zhang. Are ‘class a’ temperature requirements realistic or desirable? *Building and Environment*, 45(1):4–10, 2010.
- F. S. P. Bauman and E. A. Arens. Task/ambient conditioning systems: engineering and application guidelines. Center for Environmental Design Research, University of California, Berkeley, 1996.
- G. Brager, H. Zhang, and E. Arens. Evolving opportunities for providing thermal comfort. *Building Research & Information*, 43(3):274–287, 2015.
- R. De Dear, T. Akimoto, E. Arens, G. Brager, C. Candido, K. Cheong, B. Li, N. Nishihara, S. Sekhar, S. Tanabe, et al. Progress in thermal comfort research over the last twenty years. *Indoor air*, 23(6):442–461, 2013.

- V. L. Erickson and A. E. Cerpa. Thermovote: participatory sensing for efficient building hvac conditioning. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 9–16. ACM, 2012.
- D. Faulkner, W. Fisk, D. Sullivan, and D. Wyon. Ventilation efficiencies of task/ambient conditioning systems with desk mounted air supplies. LBNL-42597. <http://tinyurl.com/task-ambient>, 1999.
- L. R. Glicksman and S. Taub. Thermal and behavioral modeling of occupant-controlled heating, ventilating and air conditioning systems. *Energy and Buildings*, 25(3):243–249, 1997.
- T. Hoyt, E. Arens, and H. Zhang. Extending air temperature setpoints: Simulated energy savings and design considerations for new and retrofit buildings. *Building and Environment*, 88:89–96, 2015.
- C. Huizenga, S. Abbaszadeh, L. Zagreus, and E. A. Arens. Air quality and thermal comfort in office buildings: results of a large indoor environmental quality survey. Center for the Built Environment, 2006.
- A. Krioukov and D. Culler. Personal building controls. In *Proceedings of the 11th international conference on Information Processing in Sensor Networks*, pages 157–158. ACM, 2012.
- M. Paciuk. The role of personal control of the environment in thermal comfort and satisfaction at the workplace. PhD thesis, University of Wisconsin-Milwaukee, 1989.
- W. Pasut, H. Zhang, E. Arens, and Y. Zhai. Energy-efficient comfort with a heated/cooled chair: Results from human subject tests. *Building and Environment*, 84:10–21, 2015.
- L. P´erez-Lombard, J. Ortiz, and C. Pout. A review on buildings energy consumption information. *Energy and buildings*, 40(3):394–398, 2008.
- H. Zhang, E. Arens, D. Kim, E. Buchberger, F. Bauman, and C. Huizenga. Comfort, perceived air quality, and work performance in a low-power task–ambient conditioning system. *Building and Environment*, 45(1):29–39, 2010.
- H. Zhang, E. Arens, and W. Pasut. Air temperature thresholds for indoor comfort and perceived air quality. *Building Research & Information*, 39(2):134–144, 2011.
- H. Zhang, E. Arens, M. Taub, D. Dickerhoff, F. Bauman, M. Fountain, W. Pasut, D. Fannon, Y. Zhai, and M. Pigman. Using footwarmers in offices for thermal comfort and energy savings. *Energy and Buildings*, 2015.
- H. Zhang, E. Arens, and Y. Zhai. A review of the corrective power of personal comfort systems in non-neutral ambient environments. *Building and Environment*, 91:15–41, 2015.