

Personalized HVAC Control System

Mark Feldmeier

MIT Media Laboratory, Responsive Environments Group
75 Amherst St., Cambridge MA, 02142
Email: geppetto@media.mit.edu

Joseph A. Paradiso

MIT Media Laboratory, Responsive Environments Group
75 Amherst St., Cambridge MA, 02142
Email: joep@media.mit.edu

Abstract—We present a novel method of building comfort control, focused around the occupant. Custom sensing, communication, and actuation hardware were developed to locate users in a building, and measure various parameters directly on the body. These signals were used to infer user comfort and control the air-conditioning system to direct air flow where it was needed, when it was needed. A three month study of the system was conducted, with four weeks of this experimental control strategy compared to the previous four weeks of standard control. An improvement in both comfort and energy usage are shown as a result of this user-centric control system.

I. INTRODUCTION

Creating an appropriate indoor climate is essential to worker productivity [1] and personal happiness. It is also an area of large expenditure for building owners [2]. The largest consumer of energy in the United States is buildings, with residential stock accounting for 21% and commercial stock accounting for 18%, combining to 39% [3]. Within buildings themselves, the largest energy sinks are the heating, ventilation and air-conditioning (HVAC) systems. In residential applications, HVAC accounts for 26.1% of the total energy usage [4], whereas in commercial applications, HVAC accounts for 53.4% [5]. This makes building support systems, especially in the case of commercial buildings, a prime target for energy savings.

But, what can be done to reduce costs in these areas? Either more efficient ventilation technologies can be developed, or the existing technologies can be used more efficiently. Considering the long life span of buildings, and the fact that most commercial buildings are more than 15 years old [6], the latter proposition seems more cost effective, as merely adding a new control and sensing layer would be far less expensive than replacing a whole ventilation system. This idea is promoted further by the notion that most buildings are currently being run inefficiently due to the non-adaptable nature of their control systems, and that savings of up to 35% are possible [7]. As Vastamäki et al. clearly describe in their analysis of thermostat usage [8], the fundamental efficiency of the building and the comfort of the occupants both suffer when the occupant does not understand the behavior of the building. Users are shown to consistently over-turn thermostat dials in response to uncomfortable conditions, causing thermal oscillations that waste energy and create an uncomfortable environment.

Work to create such responsive environments began in earnest with the beginning of ubiquitous computing in the late

1980s. At Xerox PARC, offices were equipped with radio-frequency identification (RFID) and light, temperature, and occupancy sensors, which were allowed to turn off outlets, adjust HVAC systems, and control lighting [9]. Portable devices allowed users to edit preferences wherever they were in the building. Lee et al. [10] use a similar approach to resolve the conflicting comfort needs of users.

These very programmatic responses were challenged by Mozer, whose neural networked house [11] would purposefully turn lights off in order to understand if they were needed. Although this creates a longer learning curve than preassigned knowledge, it is capable of adapting over time without intentional user input. Adaption has also been explored by modeling buildings as multi-agent systems [12].

Ultimately, the majority of HVAC control work is focused on energy savings and temperature regulation, not human comfort. Although the control algorithms and adaptive strategies are directly applicable, the determination of personal comfort is not a solved problem. Multiple factors have been studied in their relationship to comfort, with the Predicted Mean Vote (PMV) [13] being the most common metric. The PMV averages user comfort over large populations considering temperature, humidity, wind speed, thermal radiation, activity, and clothing. This works well in practice [14], but does not fit all needs. A variety of other factors influence comfort, including age [15], local climate and culture, and the availability of natural ventilation [16].

The major use of the PMV is to set boundaries on temperature, humidity, and wind speed to a comfortable level within a building. Distributed sensor networks have been employed in attempts to assess comfort by measuring PMV values in real-time [17] [18] [19], but these involve cumbersome hardware, invasive systems, or have limited accuracy.

These previous works, with the exception of [19], attempt to assess the PMV as a global variable: a fixed standard for all people. Megri et al. [20] use PMV sensors similar to [18], and poll the user as [19] does, but instead use a support vector machine algorithm to determine the indices of the occupant's comfort, rather than using a PMV table. They show 99% accuracy at predicting comfort in this manner, which points to the possibilities of automatic recognition of comfort, on a person-by-person basis. Unfortunately this work involved large, tethered sensors which are required in three locations in close proximity to the user.

The problem is not only one of assessing an individual's

personal comfort level - an effective control system must also be able to locate the person and affect that proximate temperature. As temperatures can vary greatly, even within the same room, this proves a difficult task. Forced air distribution systems, particularly underfloor methods, can be tapped at points along the run to allow air to circulate locally. Many companies make systems [21] which implement this, usually under the name Task Ambient Conditioning (TAC). These systems allow the user to adjust air flow, and sometimes temperature, at a local vent. Not only is the availability of fresh air shown to give greater comfort, it is also more efficient [22], as air is only chilled where needed, and larger sections of the building can be allowed to drift out of normal comfort zones. TAC systems are very expensive and difficult to install after initial construction, however.

Our work addresses these problems by creating and testing at scale a uniquely adaptive control structure based upon both preassigned and dynamic knowledge. The user is only required to press a button indicating the direction of discomfort, if and when she is uncomfortable, and the building deals with the difficult tasks of dynamic energy management and conflict resolution. Specific sensor hardware is developed, which make the task of installing and operating these systems easier. Adaptive pattern recognition and control algorithms are presented, along with their efficacy in increasing the personal comfort of building occupants while reducing building energy consumption.

II. SYSTEM OVERVIEW

The Personalized HVAC System consists of four main components: portable nodes, room nodes, control nodes, and a central network hub. A diagram of a typical installation can be seen in Figure 1. These nodes work together to track user location and comfort, make control decisions, and actuate the necessary airflow sources. The portable nodes and control nodes communicate via wireless to the room nodes, which relay the data and control commands to the central network hub via ethernet. The wireless protocol is a low-power, 2.4GHz communication scheme based on the 802.15.4 MAC layer. We




Fig. 1. Typical installation with fixed and wearable portable nodes(1 – 3), room node(4), control nodes(5 – 6), and the central network hub(7).

briefly overview the system design in this section - Ref. [23] provides much more detail.

This sensor network was deployed on the third floor of the MIT Media Laboratory's Wiesner Building with four offices and one large common space being outfitted, encompassing the workspace of 10 people. Two operable windows exist in this space that were equipped with control nodes and motorized openers. There were also seven variable-air-volume (VAV) dampers, which were retrofit with control nodes and motors for regulating chilled air (no heat sources were used).

At the heart of the system is the building occupant; this is where the comfort information resides. To best assess the occupant's comfort level, a portable node was developed. This sensor node is lightweight and small enough to remain almost unnoticed by its user. It weighs 30g, and measures 54mm by 40mm by 14mm when packaged (see Figure 2). It also has low average power consumption ($11\mu\text{A}$, giving a two year coin-cell battery life), as frequent battery recharges or large battery size are nuisances to the user.

This portable node senses the local temperature, humidity, light level, and inertial activity level of the user. The continuous inertial sensing is accomplished at a fraction of what's typically required via the use of a specially designed passive piezo accelerometer with an analog activity integrator. It also has three buttons on the side, which allow the user to input current comfort state (one button each for hot, cold, and neutral). These data are sent wirelessly, at one minute intervals, to the central network hub via the room nodes. Fixed versions of portable nodes were also placed in each room, next to the preexisting thermostat, and on the exterior of the building to gather local climate data.

The actuation of the various air sources (windows and air-conditioning dampers) is achieved via control nodes (see Figure 3), which are tethered to a 24V_{DC} power source, and have a motor that opens or closes the associated mechanical element via wireless commands. They also monitor local wind speed, temperature, humidity, and light level. These data, along with the current state of the motor, are sent off wirelessly at one minute intervals to the central network hub via the room nodes.

The backbone of the system is comprised of room nodes (see Figure 4), which receive data from the portable and control nodes and act as network coordinators, ensuring that each proximate node is only talking to one device. Since each room has at least one room node, the room-scale locations



Fig. 2. Portable node as worn on a user's wrist.



Fig. 3. Control node circuit board.



Fig. 4. Room node circuit board.

of the portable nodes can be inferred from the received signal strength indicator (RSSI) of the RF device. They also assess the local temperature, humidity, light level, and passive infrared (PIR) activity level and send these data to the central network hub at thirty second intervals. Communication to the central network hub is accomplished via an on-board ethernet module.

The central network hub is a computer that receives all of the data over ethernet and processes it according to the comfort and control algorithms. It checks for current activity in the network, and only responds to new data, resetting room nodes if they are no longer active. It also backs up current system state in case of failure, and timestamps and logs each data point for offline system analysis.

III. CONTROL ALGORITHMS

To effectively handle the complex mapping requirements of this MIMO network, a hybridized control system is employed. In a hybrid system, individual nodes exchange information and trade off responsibilities in an ad-hoc, but hierarchical fashion. This fits particularly well with the topology of sensor networks, as the control layer matches the physical instantiation. A full high-level dataflow chart can be seen in Figure 5. Our system consists of seven types of software modules: Control Modules, Location Modules, Window Modules, Outdoor Modules, Thermostat Modules, Portable Modules, and Room Modules. Each node in the network has a specific instantiation of a module associated with it to keep track of its own local state and needs.

The Control Modules receive setpoint commands from the Window, Thermostat, Room, and Outdoor Modules, and make




Fig. 5. Control system flowchart.

decisions as to how to appropriately control the associated damper and window motors to reach these setpoints. They check for full open or full closed actuators, and implement a positional-integral (PI) control scheme around a small dead-band, maintaining temperature or comfort profiles without an excess of wireless communication or motor motion.

The Location Modules aggregate all the wireless transmissions from the portable nodes, and create a location table that the Room Modules can access to find out who is where. The location is based upon the strongest RSSI from the room nodes, and a voting algorithm over the past three responses is used to smooth out the data. Although location accuracy could be improved by also using the RSSI from the control nodes in the same room with a more sophisticated estimator, this simple method worked well enough for our application because of the isolating effect of metal studs in the walls, as can be seen in Figure 6 for a typical occupant over a one hour time period. Any brief confusion, as seen at the end of the plot, was generally rectified by the voting algorithm.

The Window Modules receive information from the window control nodes, and keep a log of when the window was last manually operated, which is accessed by the Outdoor Modules when determining whether to open or close the windows. The Window Modules also monitor the air speed coming in through the window, and close the windows if it becomes too windy.

The Outdoor modules receive outdoor temperature and hu-




Fig. 6. Example RSSI values across all room nodes for an occupant of Room 44.

midity information, and poll the indoor Thermostat Modules to decide whether to open the windows and let in cool outside air. In this way, the air-conditioning can be shut off to minimize energy consumption, while still maintaining a comfortable room temperature.

The Thermostat Modules track individual room air temperatures, and poll the Room Modules to determine whether or not they should be performing control actions. Based upon the Room Modules' states, the Thermostat Modules will either run normal control, setback control, or no control at all. 'Normal' control, performed when the motion detector indicates that the room is occupied but no wearable module is detected, uses the wall-mounted sensor to regulate temperature to a predetermined setpoint. In setback control, when the room is assumed to be unoccupied, the temperature is raised 4°F to conserve energy. When wearable sensors are present, the Room Modules are active, and the VAV damper is regulated by the sum of all occupants' comfort states, as inferred by the wearable nodes and portable modules. The Thermostat Modules also determine if the room is too cold, at which point the Window Modules can not open the windows (outside air was only used for cooling).

The Portable Modules keep track of each individual user, i.e. whether they are active and comfortable. The activity recognition algorithm takes into account the user's local temperature, activity mean, and activity variance to determine whether the portable node is being worn, or is just being jostled on a table. This is an important criterion, as the temperature data from a node that is not being worn can not be allowed to control the room's air-conditioning system. This problem is made more difficult by the fact that the temperature sensor has a considerable time lag, and users are often sitting quite still while working at their desks. An example of this algorithm in practice can be seen in Figure 7. Short false positive spikes are common due to the sensitive nature of the algorithm, as fast




Fig. 7. Example activity algorithm results over four days - the node was not worn on the first and last day.

system response time was desired - these have only limited effect on system performance, however.

The Room Modules poll the Portable Modules and Location Modules to find out if users are currently active, where they are located, and if they are comfortable. Based upon these findings, they either relinquish control to the Thermostat Modules, or work with all sensor data to minimize the discomfort in the rooms. The PIR motion sensor data is used by the Room Modules to know whether the room is occupied, and predict when it will next be occupied. This prediction is based upon the arrival times of the past week, and is critical to maintain the room at an appropriate temperature when users first enter. It is also inefficient to excessively cycle the air-conditioning system, so a three hour timeout is placed on room vacancy detection.

IV. COMFORT ALGORITHM

In order for the system to function effectively, it must have a metric by which to judge how far the user is from being comfortable. This places a number of difficult constraints on what types of algorithms can be used, but is required due to the fact that the output of the implemented algorithm must drive a control loop. For this reason, it must have a monotonic structure to avoid instabilities and local minima. These could be accounted for through a non-linear control scheme, but this is avoided in order to minimize the number of variables involved in determining system stability and performance.

Many standard pattern recognition techniques are inadequate for this control task, as they seek to draw boundaries around similarly labeled data, giving accurate classification, but no distinction of levels within those classes. A simple Bayesian analysis could give a probability of comfort, but would require a much more accurate model than is currently available, given the limited set of on-body comfort indices used. Only temperature and humidity are measured on-body, whereas the PMV also requires clothing level, metabolic rate, and air flow. This problem is compounded by the limited labeling of the acquired data. In comparison to the PMV's seven point scale, users of this system only have three choices: hot, cold, or neutral. This means there is no way of knowing exactly how hot or cold they are at the instant a button is pressed. This metric must be inferred from the distribution of received data.

Not only are the labeled data points ambiguous as to their level of discomfort, but they are also very sparse in their occurrence. Users are not required to press any buttons, but only asked to do so if they feel uncomfortable, limiting the amount of labeled data points to an average of about one per person per day. Another issue faced in this reduced data set is the lack of an even distribution of hot and cold labels. For some users, the room was never cold enough to make them feel cold, so only hot data points exist, giving no information by which to determine a lower limit on comfort.

For any system to be effective, it must deliver predictable results for users, or they might respond in ways which counter both the goals of the system and themselves. For this reason, a linear discriminant was chosen, which creates a clearly defined boundary from which a positive or negative distance can be measured. The Fisher Discriminant seeks the most effective rotation matrix for the given data set to produce a projection on a lower dimensional space with high class separation. It takes a statistical approach of finding the greatest between-class scatter for the lowest within-class scatter. For this case, it is a simple matter of reducing a two dimensional space (temperature and humidity) to one, with the only difficulty being in choosing a decision boundary. Rather than the usual approach of using the intersection of sample distributions, the decision was based upon discriminating between points that represent the boundary conditions. In this case, the most representative training points are assumed to be those with the most extreme values (e.g. 'hot' data with the lowest temperature values), and a separating line is created at the mean of these data. The comfort distance can then be simply computed as the distance to this boundary.

In order to accommodate the adaptation of the comfort algorithm, a limited set of data points were used in creating the decision boundary, with new points replacing the old. Nine points each of 'hot' and 'cold' labeled data were used, which allowed a complete update in two to three weeks (users press buttons on average once a day), which was enough time for users to adapt the system before the end of the experiment. In cases where nine data points were not available, as many as were present were used. If less than two data points existed,




Fig. 8. Example decision boundary for Fisher Discriminant.

two points were selected that created a reasonable line in comparison to other users. A representative result of the Fisher Discriminant is shown in Figure 8, and more details on the analysis are given in [23].

The distance between the 'hot' and 'cold' labeled points varied greatly for the different users. Accordingly, the calculated comfort distance also varied greatly between users. To normalize this reported comfort distance so the control system could effectively arbitrate between users, the mean distance of 'hot' and 'cold' points from the decision boundary was calculated. As new temperature and humidity data were collected by a user's portable node, the final comfort output was computed as the comfort distance divided by this mean distance. In this way, it is assumed all users are equally uncomfortable for a given comfort value.

V. EVALUATION

In order to assess the efficacy of a body-worn comfort control system, a long-term user study was performed from May 18th through August 21st (summer) of 2009. The study was carried out with a mostly graduate student population at the MIT Media Laboratory. Ten people were assigned individual portable nodes, and four offices and a large common area were equipped with room nodes and control nodes.

Phase One of the study ran from May 18th to June 21st. No actuation was performed for this period, and data were merely gathered on how users interacted with the devices. In order to have a fair baseline for comparison, the maintenance department made repairs to the VAV damper controllers and thermostats for all of the offices and common spaces, certifying that the basic 1985-vintage HVAC operation met their standard. Phase Two of the study ran from June 22nd to July 26th. This stage consisted of baseline comfort data collection. Users were asked to complete a questionnaire regarding their comfort under the current HVAC system. They were also asked to wear the portable nodes, and press a button whenever they felt inclined to do so (either hot, cold, or neutral).

Phase Three ran from July 27th until August 21st. During this period, the experimental control system was run, with the HVAC system and window motors being controlled via the various sensors in the network as described earlier. Users were told that the climate control system would respond to their wearable sensor, and would try to mediate the comfort preferences of each occupant of the individual office or common space. Periodic surveys were administered during this period to assess the users' comfort level, and their beliefs about the system. Phase Three is subdivided into three sections: Week Two, Week Three, and Week Four.

A. Energy Metrics

The only method the system had available to measure energy usage was via the air flow sensors retrofitted onto the VAV boxes. Although energy monitors on the chilled water lines and fan motors would give more accurate results, the area being controlled by the experimental system was a small percentage of the total space being cooled, hence any effects would be unnoticeable. Nonetheless, our sensors show the fan usage for the space very accurately. Since fan energy usually represents 40% of total HVAC power consumption [24], this is an important metric by itself. The chiller energy can be estimated by multiplying by the number of cooling days. Cooling days are an integration, over an entire day, of the outdoor temperature difference from 65°F. This is often used to determine how much energy is required to cool a space, as it represents the temperature difference the HVAC system must produce, and has been shown to give linear correlations [25]. Our fan data is thus divided through by the number of degree cooling days to make the results more generally applicable - this energy metric, showing each room's contribution, can be seen in Figure 9.

It can be clearly seen that the total chilled air used decreased under active control. A cursory approximation of energy savings, comparing Phase Two and an average of Week Three and Week Four, shows a reduction of 75%. This is based upon an estimated 40% fan usage times the normalized VAV air usage, plus an estimated 60% chilled water usage times the non-normalized VAV air usage. The actual savings are much smaller for two reasons. Firstly, the main savings shown are due to Room 36 and Room 40 reductions, which represent an unfair comparison, as the area they cooled was also serviced by eight other VAV boxes, none of which were under experimental control. Secondly, the HVAC system, despite having been repaired, was not running perfectly for all rooms during Phase Two.

Just looking at energy consumed by the ventilation system, we saw an 8% average decrease in chilled air usage per degree cooling day, which, when multiplied by the 40% fan usage metric, can be taken as a lower energy savings bound of 3.2%. This, however, doesn't include reduction in refrigeration, which will contribute very significantly to energy savings. As we couldn't measure this directly in these tests, we can infer it from temperature differences. To best account for this energy usage, Room 44 will be analyzed




Fig. 9. Chilled air usage per cooling degree day by room.

in detail. This room is selected because it was functioning perfectly during Phase 2, having had both the thermostat and damper motor controller replaced. During Phase 2, its average regulated temperature was 73.1°F. Comparing this to Week Three's data, as this week had a very similar average number of cooling degree days to the baseline, gives an increase of average temperature to 73.5°F, a decrease in air usage of 17%, and a decrease in air usage per cooling degree day of 4%, summing to an estimated 11.8% energy savings. The average temperature increase across all rooms was 0.8°F. And, assuming a linear correlation between temperature change and energy savings, this average temperature increase would indicate a 24% energy saving. Additional insight into energy saving is provided in [23].

B. Comfort Metrics

There are two ways in which the comfort of the experimental subjects was measured: through 'hot' and 'cold' button presses and through weekly surveys. A comparison between the Entrance Surveys and Exit Surveys is shown in Figure 10, with the Entrance Survey representing user beliefs under normal unretrofit HVAC control, and the Exit Survey referencing the four weeks of experimental control. These surveys had identical questions, and were taken two months apart from each other, making them a relatively unbiased indicator of user preferences.

During Phase Three, weekly polls of thermal comfort level were performed. These employed the seven point scale used in the PMV, and can therefore be compared to standard HVAC practices of keeping the temperature within bounds of 80% occupant satisfaction. The PMV counts the 'Slightly Cool' through 'Slightly Warm' categories as being comfortable, and the occupants were in this zone 81% of the time. This percentage increased over the study, starting at 76% and ending at 85%, most likely due to the system learning the preferences of the users. Details are presented in [23].

For the majority of the users, over 95% comfort rates were common. In fact, the only users for which this is not the case are those for which 'cold' data points had to be manually entered (as they never punched 'too cold' on their wearable),




Fig. 10. Response to entrance and exit surveys: “Overall, I am satisfied with the comfort level in my office”.

and User 20, whose data was not received during Week Two because of radio problems. User 20’s comfort increased from 20% to 75% after this bug was fixed, suggesting that our results are not merely placebo gains.

An insight into how well our comfort calculation algorithm performs is given in Figure 11, which shows a plot of users’ reported comfort level versus the system’s beliefs of their comfort. The data is taken from weekly polling of users, ranging from -3 (‘cold’) to $+3$ (‘hot’). The computed comfort uses a similar metric, with greater positive values representing greater heat discomfort. Nodes that were partly malfunctioning at the time this data was taken are shown with asterisks in the plot, to help differentiate their performance from the functioning nodes (sensor data, hence comfort level, is still valid here, but the control algorithm wasn’t functioning properly, which resulted in the higher discomfort levels). Lines are also drawn on this plot to indicate the deadband within which the active comfort control algorithm would decide that no action was required.




Fig. 11. Reported comfort versus computed comfort for all users: Week Two through Week Four.

With the exception of a few points, the computed comfort tends to follow the reported comfort. It is also important to note that the level of reported comfort used in the PMV to determine comfort boundaries is ± 1 , which is where all the functioning nodes lie, indicating that our control system is working properly. Even for the majority of the malfunctioning nodes, the system knew that those users were hot, but was unable to respond accordingly.

This personalized system also arbitrates between users sharing a common space. A plot of two users’ computed comfort distances can be seen in Figure 12, along with the associated VAV damper air flow, to indicate the control system’s actions. When User B enters the room, it can be seen that the damper shuts off the airflow, as that user is assumed to be cold. When User A enters at 2h, and is assumed hot because of her different comfort profile, the system increases the air flow in an attempt to average the comfort of the two users. The damper is full open, and the comfort distances change only slightly, although in the correct directions. When each user leaves, the system again compensates correctly, as quickly as the system response and room heat capacity allow, which is often quite slow.

Despite this fairly long response time, users still felt that the system was responding to their needs, and trying to average between all occupants. 80% of the users responded favorably to the survey question, “I believe the personalized comfort system is doing a good job of balancing the thermal comfort needs of all the people in my workspace.” From personal recollection, users would often look forward to their officemates leaving, so they could have a more comfortable environment.

Another metric of personal comfort is the number of button presses made during the experiment, since each user had the ability to register his or her personal comfort at any point in time via the wearable nodes. As detailed in [23], the average number of button presses decreased 14% between Phase Two




Fig. 12. Example arbitration between users.

and Phase Three. This is not a dramatic change, but it is significant, due to the fact that there was greater incentive during Phase Three to press buttons, as this would actually change environmental conditions. Users were clearly being made more comfortable by the experimental system, despite slow control responses.

VI. CONCLUSIONS

Although the energy savings vary depending upon assumptions, it is clear that our experimental control system significantly reduced chilled air usage, and we estimate an energy savings of up to 24% over the standard HVAC control system that was running previously. It accomplished this by only cooling areas as much as required to maintain occupant comfort, and not cooling areas when occupants were not present. It also worked to maintain room temperatures at an equitable level for all involved. It was able to do all of this as a result of an ultra-low-power, wrist worn sensor node, which made the building aware of its occupants' state via sensor data, and simple 'hot' and 'cold' button presses.

These energy savings were the direct result of improving user comfort. In a well-functioning building, it is not the case that the temperature is either too hot or too cold, but rather that it is too hot or too cold for particular individuals: the air is not being distributed effectively. A number of personal cooling systems are commercially available, but these are expensive, and in some cases impossible to install (e.g. underfloor systems). Our system overcomes these limitations by using wireless sensors and actuators, enabling the retrofit of older and less efficient buildings. Indeed, it exceeded the 80% comfort goal, which all commercial buildings aspire to, but very few meet.

Our future work will explore longer studies involving more people and an entire building that we can control year-round without extensive retrofitting (as it is already equipped with a modern HVAC system). We will also explore variations on the wearable sensor, perhaps wearing it elsewhere on the body or integrating it into a watch, which would be more comfortable for our users. We will explore ways of better accommodating transients (when users transition between thermal environments), and mining all sensor data (including integrated activity level) for inferring comfort. We are also integrating our system with mobile phone applications and a pervasive display network installed throughout our building to explore aspects of persuasive computing - showing users how the system is working to keep them comfortable while reducing energy consumption and encouraging them to accept a wider range of comfort in exchange for energy impact.

ACKNOWLEDGMENTS

This work was partially funded by an anonymous donor through the MITEI, a fellowship from the Martin Family Society of Fellows for Sustainability, the MIT Media Laboratory, and the Things That Think consortium.

REFERENCES

- [1] W. Fisk and A. Rosenfeld, "Estimates of improved productivity and health from better indoor environments," *Indoor Air*, vol. 7, pp. 158–172, 1997.
- [2] J. Kreider, P. Curtiss, and A. Rabi, *Heating and Cooling of Buildings: Design for Efficiency*. McGraw-Hill, 2002.
- [3] U. E. I. Administration, "Energy consumption by sector," *2006 Annual Energy Review*, 2007.
- [4] U. E. I. Agency, "Residential end use electricity consumption," *Residential End Use Consumption Survey, 2001*, 2007.
- [5] U. E. I. Administration, "Preliminary end use consumption estimates," *Commercial Building End Use Consumption Survey, 1999*, 2003.
- [6] R. Diamond, "An overview of the US Building Stock," *Indoor air quality handbook*, 2001.
- [7] Federspiel Controls, <http://www.federspielcontrols.com>.
- [8] R. Vastamäki, I. Sinkkonen, and C. Leinonen, "A behavioural model of temperature controller usage and energy saving," *Personal and Ubiquitous Computing*, vol. 9, no. 4, pp. 250–259, 2005.
- [9] S. Elrod, G. Hall, R. Costanza, M. Dixon, and J. Des Rivieres, "The responsive environment: Using ubiquitous computing for office comfort and energy management," Technical Report CSL-93-5, Xerox Palo Alto Research Center, Tech. Rep., 1993.
- [10] H. Lee, J. Choi, and R. Elmasri, "A Conflict Resolution Architecture for the Comfort of Occupants in Intelligent Office," in *Intelligent Environments, 2008 IET 4th International Conference on*, 2008, pp. 1–8.
- [11] M. Mozer, "The neural network house: An environment that adapts to its inhabitants," in *Proceedings of the American Association for Artificial Intelligence Spring Symposium on Intelligent Environments*. Menlo Park, CA: AAAI Press, 1998, pp. 110–114.
- [12] U. Rutishauser, J. Joller, and R. Douglas, "Control and learning of ambience by an intelligent building," *IEEE Transactions on Systems, Man and Cybernetics, Part A*, vol. 35, no. 1, pp. 121–132, 2005.
- [13] P. Fanger, *Thermal Comfort*. Copenhagen: Danish Technical Press, 1970.
- [14] R. De Dear, "Thermal comfort in practice," *Indoor Air*, vol. 14, no. s7, pp. 32–39, 2004.
- [15] N. Taylor, N. Allsopp, and D. Parkes, "Preferred room temperature of young vs aged males: the influence of thermal sensation, thermal comfort, and affect," *Journals of Gerontology Series A: Biological and Medical Sciences*, vol. 50, no. 4, pp. 216–221, 1995.
- [16] J. Nicol and M. Humphreys, "Adaptive thermal comfort and sustainable thermal standards for buildings," *Energy & Buildings*, vol. 34, no. 6, pp. 563–572, 2002.
- [17] B. Ivanov, O. Zhelondz, L. Borodulkin, and H. Ruser, "Distributed smart sensor system for indoor climate monitoring," in *KONNEX Scientific Conference, München*, 2002.
- [18] W. Tse and W. Chan, "A distributed sensor network for measurement of human thermal comfort feelings," *Sensors & Actuators: A. Physical*, vol. 144, no. 2, pp. 394–402, 2008.
- [19] C. Federspiel and H. Asada, "User-adaptable comfort control for HVAC systems," in *Journal of Dynamic Systems, Measurement, and Control*, vol. 116, Sep 1994, pp. 474–487.
- [20] A. Megri, I. Naqa, and F. Haghghat, "A Learning Machine Approach for Predicting Thermal Comfort Indices," *International journal of ventilation*, vol. 3, no. 4, pp. 363–376, 2005.
- [21] F. Bauman, A. Baughman, G. Carter, and E. Arens, "A field study of PEM (Personal Environmental Module) performance in Bank of America's San Francisco office buildings," *Berkeley, Center for Environmental Design Research, University of California. Final report submitted to Johnson Controls World Services, Inc*, 1997.
- [22] D. Wang, E. Arens, T. Webster, and M. Shi, "How the number and placement of sensors controlling room air distribution systems affect energy use and comfort," in *International Conference for Enhanced Building Operations*, 2002.
- [23] M. Feldmeier, "Personalized building comfort control," Ph.D. dissertation, MIT, 2009.
- [24] D. Westphalen and S. Koszalinski, "Energy Consumption Characteristics of Commercial Building HVAC Systems. Volume II: Thermal Distribution, Auxiliary Equipment, and Ventilation," 1999.
- [25] D. Claridge, J. Haberl, S. Katipamula, D. O'Neal, D. Ruch, L. Chen, T. Heneghan, S. Hinchey, J. Kissock, and J. Wang, "Analysis of Texas LoanSTAR Data," in *7th Annual Symposium on Improving Building Systems in Hot and Humid Climates*, 1990, pp. 9–10.