ThermoNet: Wireless Solution for Fine-Grain Building Comfort-Efficiency Assessment

Jing Li, Jin He, and Anish Arora

Department of Computer Science & Engineering
The Ohio State University
Columbus, OH 43210
{ljing, heji, anish}@cse.ohio-state.edu

***Abstract*—Understanding effectiveness of the Heating, Ventilation, and Air Conditioning (HVAC) system in a building is a prerequisite for energy efficiency optimization. This particular context has, to our knowledge, not yet received adequate attention thus far. To address this issue, we evaluate the thermal comfort and the energy efficiency of a relatively modern HVAC in a large building based on building wide high-fidelity environmental data collected over 12 months. Access to fine grain information reveals temporal and spatial dynamics that help quantify the level of (non-)compliance with control objective and green-build thermal comfort standards, over-conditioning at multiple time scales that offer opportunities for reduced operating cost, and identify not only persistently ill-conditioned rooms but also intermittently ill-conditioned rooms that need maintenance. The paper moreover describes ThermoNet, our hybrid wireless sensor network solution for monitoring a large building, which uses duty cycling and adaptive power control to achieve high data yield with low power consumption.**

***Keywords*-HVAC System; Building; Thermal Comfort; Energy Efficiency; Wireless Sensor Network**

Figure 1. Dreese Laboratory in the main campus of The Ohio State University.

I. Introduction

Given that approximately 35% of the energy in the United States was used for Heating, Ventilation, and Air-Conditioning system [1], a number of efforts have been made over the past decade to improve HVAC performance by instrumenting Wireless Sensor Networks (WSNs) in residences, campus offices, and production data centers [2], [3], [4], [5], [6], [7], [8], [9]. A major theme of these research has been leveraging occupancy information to optimize HVAC schedules or adapt setbacks, where either per-room controller is considered or uniform effectiveness of a building wide HVAC system is assumed. However, thermal and air dynamics in large spaces tend to be complex, especially in modern buildings accommodating hundreds if not thousands of rooms. This is due to interdependencies in the design of air distribution/diffusion, customization of air handling to the respective functions of diverse rooms, and selection of HVAC control policies. Thus, a fine grain instrumentation and evaluation of an actual HVAC system provides not only realtime feedback and alarms to controller but also long-term statistics of the environment for optimizing building control schedules.

Installing, wall-powering, and wiring instruments (such as temperature, humidity, light sensors) in every room of a large building involve a significant amount of work and expenses, especially for an already constructed building since interfacing with existing building management system is cumbersome and costly. To tackle this problem, this paper presents a promising low-cost WSN solution, ThermoNet, which analyzes HVAC performance in terms of thermal comfort versus energy consumption based on dense (i.e., per room), long-lived environmental monitoring of a relatively large building. Our design of ThermoNet was based on several key objectives. First, the system should be incrementally deployable without requiring costly modifications to the building. Second, we wanted the system to be as low cost as possible so as to span all floors and rooms. Finally, the system should be able to communicate reliably to every corner of the building for a period that lasts years.

Specifically, we draw conclusions about the HVAC performance of a 9-floor New Dreese Laboratory[1] at the Ohio State University, which was added to an existing structure in 1994, based on year-long data from a 100+ node ThermoNet, which reports temperature and light data from essentially every room in the building once every 15 minutes. The building, pictured in Figure 1, houses the Computer Science & Engineering Department in more than 100 rooms which serve as offices, labs, conference rooms, classrooms, and compute centers. In contrast to BACnet[2] enabled buildings, rooms in Dreese are not equipped with environmental sensors; thus, building operators do not have access to realtime or historical data from each room.

We begin with a description of the target objective and architecture of the HVAC in question. The current target is to maintain the whole building at $72^\circ F$ with ± 2 degree error during daytime all through the year, and between $60^\circ F$ and $80^\circ F$ at night. The university has recently adopted a Green Building policy compatible with the U.S. Green Building Council (USGBC) Leadership in Energy and Environmental Design (LEED) Silver-level certification, which mandates maintaining temperatures during heating and cooling periods at $70^\circ F$ and $76^\circ F$ with ± 2 degree error respectively when occupied [10]. The control policies for the HVAC in Dreese Laboratory are however yet to be modified to meet this policy. The architecture includes two standard Variable Air Volume (VAV) Air Handler Units (AHUs), which respectively reset their discharge air set points between $55^\circ F$ and $65^\circ F$, by periodically comparing a reference set point with the highest of the temperatures sampled on-line from three pre-selected rooms. The two are shut down between 11:30pm at night and 5:30am/7:00am respectively in the morning. Shown in Figure 2, a second tier of independent (pneumatic) control exists in approximately half of the rooms, whereby the room can control its discharge set point with a change in air velocity and a reheat coil. However, the local controller is not designed for building maintainers instead of room occupants, which is enclosed by a metal lid (see the left one of Figure 2). Occupants in Dreese rooms are not allowed to open or adjust the controller; thus building occupants have little impact on room temperature set points.

Thermal comfort is a complex measurement that depends on many aspects. The most common comfort measurement is Fanger's Predicted Mean Vote (PMV) as standardized in ISO 7730 [11], depending on air temperature, radiant temperature, humidity, air velocity, occupants clothing and activity. However, measuring the PMV based on this model is very complex due to the many influences, e.g., radiant temperature and air velocity sensors are expensive and complex, such that large scale installations are not affordable.

[1] In the rest of the paper, New Dreese is abbreviated to Dreese.

[2] BACnet protocol defines a number of services that are used to communicate between building devices, which allows communication of building automation and control systems for applications such as HVAC control and lighting control and their associated equipment.

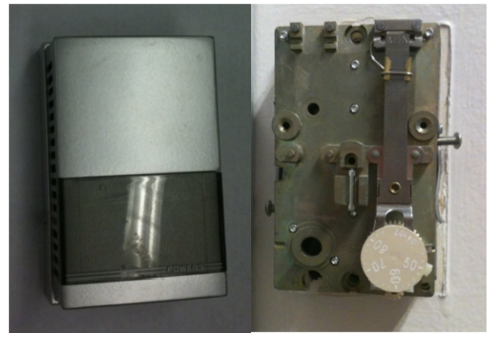

Figure 2. A pneumatic controller on the wall of a Dreese room in its enclosure (left) and exposed (right).

Given these complexities, our work focuses on evaluating temperature rather than attempting to measure PMV. Broadly speaking, as detailed in the following sections, our findings regarding the building's HVAC operation are as follows.

- In terms of thermal comfort with respect to its control objective, although the building-wide average temperature is typically within limits, the percentage of the comfortable area to the whole area is on average only 47% of the building (and it meets the Green Build Policy even less).
- Improving thermal comfort needs a refined control policy that computes the AHU discharge set point based on more than three statically chosen rooms.
- A substantial opportunity for energy savings exists based on dynamics at different time scales: (i) Many rooms are overcooled (overheated, respectively) in summer (in winter, respectively). (ii) During season changes, switches to the cooling and heating modes occasionally yield overcooling and overheating respectively. (iii) Daily patterns of building temperature and (illumination-based) occupancy also show that the AHUs can shut for several more hours at night while still meeting the control objective.
- There are not only a significant number of persistently ill-conditioned rooms, but also a significant number of intermittently ill-conditioned rooms, that warrant maintenance for thermal comfort.

II. Related Work

Research has shown that lack of visibility into building operating conditions is a root cause for low energy efficiency. Over the last decade there is a growing consensus that low cost wireless sensor networks are appropriate solutions to collecting high fidelity instruments from the environment. RACNet [5] was developed with this motivation and is perhaps the most related work to ours. Towards understanding the thermal conditions in a production data center, it developed an acquisition protocol to monitor temperature and humidity from 52 sensor nodes. Its modeling effort is more dependent on the compute activity patterns and center layout, whereas ours is more room- and function-centric in addition to depending upon building/system layout. Different investigated environment leads to discrepancy in monitoring network topology, data acquisition protocol design, and data analysis.

Traditional temperature sensors require dedicate wiring for energy supply and information retrieval [12], [13]. In recent years, there have been many emerging companies providing thermal auditing or even data analytic services so as to achieve energy conservation at homes. For instance, THUM [14] is a USB temperature/humidity sensor ($160) which connects to PC through an available USB port. The included software on PC stores all readings to a database and can be set up to email the user if the designated alarm point is reached. However, the additional costs of THUM include an always-on PC with available USB port and sensing data is only accessible by logging in remotely from any Internet-connected computer. The Venstar wireless thermostat system ($290) [15] consists of two parts: the thermostat is a wireless, battery operated unit that can be placed in any location of a home (up to 500ft. from receiver) and will transmit radio signals to the receiver which replaces the existing wall thermostat; the receiver has to be wired like a standard 24VAC thermostat, which supports up to 4 thermostats. The Venstar wireless device, which is designed for a home environment, is not compatible with existing centralized HVAC system in buildings like Dreese. In contrast, ThermoNet is developed with building environment in mind, offering non-intrusive low-cost thermal instrumentation for a large HVAC system.

On the other hand, leveraging occupancy information to optimize HVAC schedules or adapt setbacks has been a major theme of recent research. For instance, a self-programming thermostat [2] has been shown to automatically create an optimal temperature setback schedule based on occupancy statistics of a home. The smart thermostat suits best for an individual room instead of a central HVAC system controlling hundreds of rooms. At UC Merced, occupancy of multiple rooms has been estimated by a wireless camera sensor network [4]. The authors constructed models of occupancy based on Markov Chains trained with ground truth data and suggested occupancy-based ventilation control strategies accordingly. Later on they show in EnergyPlus simulator that it is possible to achieve 42% annual energy savings while still maintaining ASHRAE comfort standards by applying their predictive control algorithm [1]. These efforts assume a generic room-based HVAC system in contrast to our work, where the main focus is evaluating overall effectiveness of a specific HVAC system in a large building. Agarwal et al. in [3] presented the design and implementation of a presence sensor platform that is used for accurate occupancy detection at the level of individual offices at UCSD. Their follow up paper [16] further describes its use in the operation of HVAC resources in the CSE

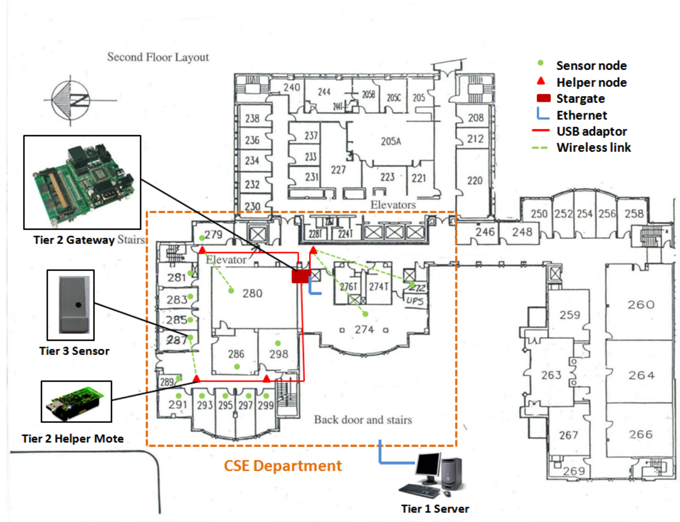

Figure 3. ThermoNet architecture illustrated for the 2nd floor of Dreese Laboratory.

building. By providing real-time occupancy information to the campus-level management network, they implemented an HVAC control scheme that obtains approximately 8% to 15% savings in both HVAC electrical and thermal energy use through controlling just a single floor of their four floor building. Their investigated building is six year old, which is controlled through a BACnet control network that provides thermal measurements to researchers as well as access to each thermal zone. In our case, Dreese Laboratory was a relatively old building constructed in 1994 when BACnet was not the ASHRAE/ANSI standard yet. Nevertheless, Dreese stands for the situation of the majority of buildings in our campus; hence, studying the effectiveness and efficiency of its HVAC system is materially useful.

III. System Design Rationale

An in-building surveillance application requires a low-cost data acquisition system that covers every corner of the building. We choose low-power sensor network for theoretical and practical considerations. Firstly, the investigated building control system does not have channel to per-room thermal conditions except three predefined sampler rooms. Secondly, interfacing with existing building management system is cumbersome. Wiring and wall-powering every sensor node is not feasible in Dreese because it will involve a significant amount of work and expenses. Thirdly, network based on IEEE 802.15.4 radios are more attractive than Wi-Fi or Bluetooth radios given that low-power radio solution leads to lower overall energy consumption and cost. Although these radios have challenges such as lower data throughput, our application only requires low data transmission rate from each source. Detailed costs of deploying ThermoNet will be elaborated in the following subsection.

A. System Architecture

ThermoNet is a three-tier sensor network, consisting of a server, a backbone, each supporting tens of wireless sensor nodes. As shown in Figure 3, Tier 1 consists of the base station, a Dell PowerEdge series server, which collects and stores sensor measurements in a database. It also maintains and supports visualization of temperature and light statistics at different time scales.

Tier 2 consists of gateway nodes which are located in a control room on the floor, and 4-5 helper motes per gateway node. This backbone enables messages from sensor nodes to be forwarded to the server, and vice versa. These gateway nodes are Stargate 7.2 devices [17], which are 32-bit class Linux device, and the helper motes are low-power TelosB motes [18]. Daemons running on Stargate forward packets received from helper motes to the base station and vice versa. In this deployment, Stargate, placed in the control room of a floor, connects to the base station via department's Intranet. Stargates are wire connected to their helper motes through USB adaptors; these wires are strung through the false ceiling of a floor, which are invisible to building occupants, with helper mote locations being chosen carefully to ensure full coverage of the floor. Field testing suggests that one floor of helper motes can communicate with sensor nodes on 3-5 floors reliably, thus a recommended backbone scale in Dreese requires three gateway nodes, e.g., on floor 2, 5, and 7. By constructing the backbone sensing fabric, Tier 3 sensors are able to send and receive packets from almost every corner of the building, enabling a variety of applications running on top of it.

Tier 3 consists of TelosB temperature and light sensor nodes, deployed in almost every room of the building (printer, control, restricted rooms, and restrooms were excluded). These nodes report the measurements made in each room periodically in wireless fashion via the Tier 2 helper nodes. These nodes are powered by a pair of AA batteries and packaged in an enclosure we designed and manufactured in China (see Figure 4), so as to be aesthetically and location-wise compatible with the control unit in each room. The cost per enclosure was about US$4; it includes small vents near the temperature sensor and has a hole close to the light sensor, to allow proper measurements. We typically placed the packaged sensor motes abutting (and lightly glued) to its room control unit, the top of white board or bookshelf when room control unit does not exist. Light sensor comes with the ordered TelosB mote and the temperature sensor is TMP36 [13], featuring low cost and low power consumption consuming less than $50\mu A$ current. It provides a voltage output that is linearly proportional to the Celsius temperature. Before deployment, we tested the accuracy of TMP36 sensors by comparing their readings with that of a high-precision thermostat. We find that the sensor accuracy is less than $\pm 2°F$ over the $32 - 122°F$ temperature range.

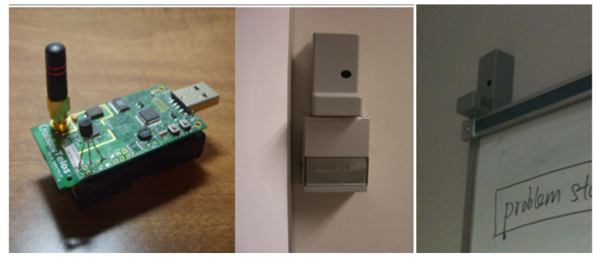

Figure 4. Sensor node, exposed (left) and in its enclosure as deployed in rooms of Dreese (middle and right).

This partly-wireless, hierarchical sensing structure was chosen for multiple reasons. First, backbone nodes enable deployment for a broad range of sensing applications (other than temperature/light monitoring). Second, the tiered topology allows potential shifting of energy and computation intensive tasks from battery-powered sensors to wall-powered backbone nodes. It also simplifies maintenance to some extent since the backbone nodes are accessible since they are in common (ceiling) areas, whereas the sensors are in rooms where we need permission from occupants to enter. Third, the architecture enables debugging and reprogramming of TelosB nodes from the Stargate.

In terms of cost, the gateway package on each floor contains one Stargate ($300), 5 ethernet adaptor ($30 each), 5 TelosB motes ($70 each in 2005), and one USB hub ($10), amounting to $810 for one floor. The per-floor wiring of helper nodes to gateway nodes took roughly 2 hours of labor, i.e., approximately $60. A room uses one TelosB mote mounted with a temperature sensor ($0.61 each), a pair of AA batteries ($1), and an enclosure ($4), leading to a cost of $76 per room. A typical base station costs $500. Hence, the amortized per-room cost of ThermoNet by deploying three floors of gateway nodes for ∼100 rooms in Dreese amounts to

$$C = \frac{\$500 + \$810 \times 3 + \$60 \times 3 + \$76 \times 100}{100} = \$107.1,$$

where the costs of deploying sensor infrastructure is amortized to approximately $37 per room. Note that in the experiment, we used TelosB motes which were designed for general WSN experimental purposes at a relatively high price. If customized sensor mote is used, the per room cost would be reduced significantly.

B. Dynamic Power Adaptation

Except for the backbone TelosB nodes which are powered through USB adaptors, ∼100 Tier 3 TelosB motes rely on AA batteries for energy supply. Their two-year lifetime is achieved by duty cycling as well as power adaptation, as follows.

First, a Tier 3 node is duty cycled at approximately 0.003%, i.e., every 15 minutes the node turns on its radio to

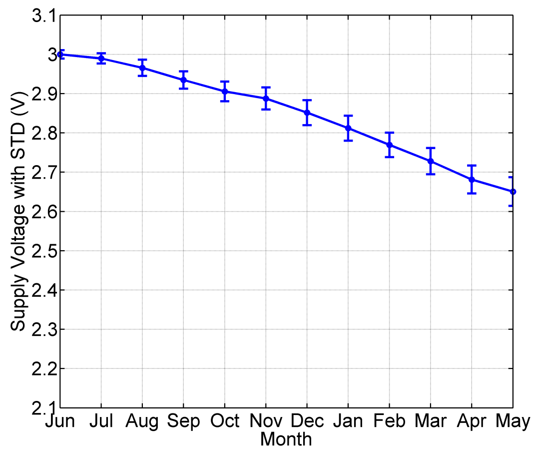

Figure 5. Average supply voltage with std at Tier 3 nodes between June 2010 and May 2011.

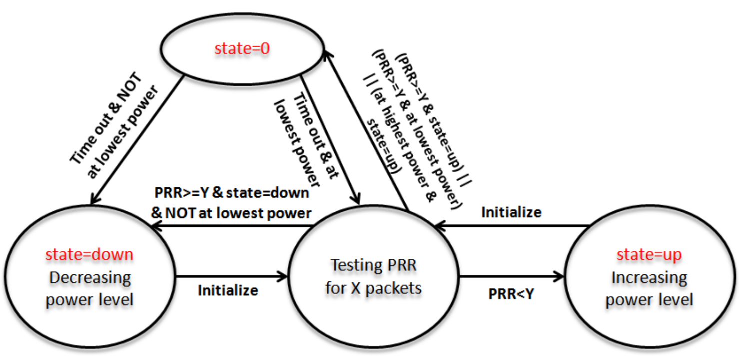

Figure 6. State diagram of power adaptation algorithm in ThermoNet.

report data, waits for an acknowledgement from backbone nodes for 20 milliseconds, and then switches off its radio. The base station keeps track of periodically sampled node health information as well as instruments, containing voltage, transmission power level, and other network relevant information. Since backbone nodes are always on, packets sent from Tier 3 will be relayed to the server with trivial delay. Figure 5 analyzes from traces the average voltage change over time on Tier 3 nodes with standard deviation. Since TelosB datasheet [18] suggests that the minimal supply voltage for mote operations is $2.1V$, the expected lifetime of a typical Tier 3 node is more than two years, longer than the theoretical expectation given that a conservative power level (the highest) was used in our calculation. As will be described in next paragraph, the power adaptation scheme further reduces energy consumption on wireless nodes. When the supply voltage drops to the lower bound, AA batteries need to be replaced at Tier 3 nodes.

Second, the power adaptation component minimizes a node's transmit power level subject to the constraint of achieving reliable communication. Since it is known that RF channels in buildings vary over time and space, in part

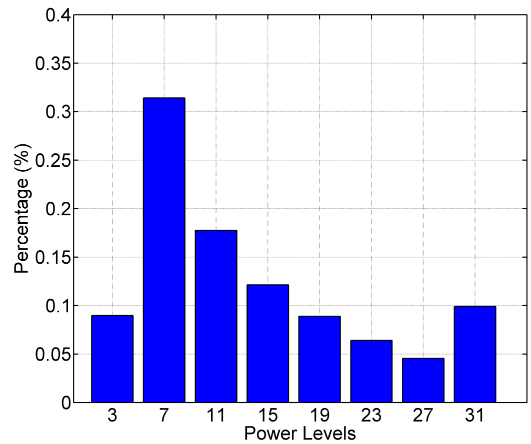

Figure 7. Probability Mass Function of transmission power level at Tier 3 nodes to achieve 98% PRR over the year.

due to the existence of a number of WiFi access points, it is thus infeasible for a node to maximize both reliability and efficiency by choosing a constant power level. Our solution for selecting power levels is illustrated in Figure 6. Initially each node starts with the highest transmission power level, i.e., 31. The power adaptation component periodically reduces the current power of a node by a unit of four levels (state=down) until the measured Packet Reception Ratio (PRR) drops below the threshold of 98%. Then, it increases power gradually to the lowest level that meets the PRR requirement (state=up). Aside from conserving energy usage, another advantage of the power reduction procedure is that Tier 3 nodes can lock to the backbone node(s) with best link quality in terms of communication cost while eliminating less reliable links. Data traces from ThermoNet reveal that a wide range of power levels have been adopted by Tier 3 in communication. We plot the Probability Mass Function (PMF) for transmission power levels used by Tier 3 nodes over a year in Figure 7, where the power range of 3 to 31 is divided into eight operating levels. It is shown that the dominant transmission power level has been reduced to the second lowest level (7) across the network.

C. Bidirectional Channels

To avoid physically accessing sensors, ThermoNet supports two-way communication between applications and sensors, i.e., in addition to retrieving data from sensors, users are capable of sending commands to change sensor configuration through the fabric. For example, users may want to change the sampling rate from a particular set of sensors for different purposes.

In order to rendezvous with duty-cycled sensors, the application-to-sensor channel is implemented by piggybacking a structure of command index and destination to the acknowledgement sent by a backbone node. Whenever a sensor node receives an acknowledgement, it checks the

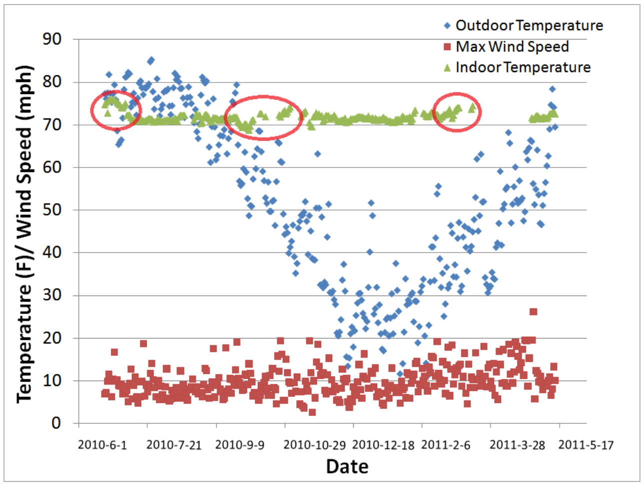

Figure 8. Long-term indoor and outdoor weather comparison.

structure by comparing the received index with its own command index, and execute the requested command if the received one is fresher (i.e., larger) and destined to it. In arbitrary states, when multiple backbone nodes carry different command indexes, the largest index carrier would succeed in disseminating its command to sensor nodes. This feature is useful in debugging Tier 3 nodes, e.g., we re-initiated the states of 2 motes when junk messages were received from these nodes at the base station.

IV. Building Thermal Comfort Assessment

This section quantifies level of (non-)compliance with the current thermal comfort objective. This helps identify opportunities for improving the HVAC system as well as identify anomalous rooms whose proper conditioning requires manual maintenance or local controller adjustment. We call these anomalous rooms *ill-conditioned*, and discuss a classifier that discriminates different types of ill-conditioned rooms. These include refining heating/cooling control policies at the global HVAC controller, adjusting local controllers for intermittently ill-conditioned rooms, and performing maintenance associated with persistently ill-conditioned rooms.

Seasonal Mode Change Anomaly. We begin by characterizing indoor daily temperature with respect to outdoor ambient conditions, such as air temperature and wind speed. Based on year-long sensing data from ThermoNet and external weather conditions measured by an on-campus weather station [19], we compare the average indoor temperature, average outdoor air temperature, and maximal wind speed in Figure 8. The indoor temperature is defined as the mean of room daily temperatures, the value of which varies in the range of $[69^{\circ}F, 76^{\circ}F]$. Note that discontinuities of indoor data points are due to periodic network maintenance performed at gateway nodes.

We find that the correlation coefficients between indoor and outdoor temperatures (indoor temperature and wind speed, respectively) over the year are within ± 0.3, implying that ambient conditions have little influence on the building. More importantly, as indicated by the ovals in Figure 8, we discover that indoor temperatures tend to deviate from the norm during certain periods of time, typically when season changes occur. For example, during the period of the second oval, the building was first overcooled and then overheated as outdoor air temperature started dropping in October. Our discussions with building automation colleagues suggest that this is likely due to HVAC malfunction, although there is some chance that if building policies were refined that these could be avoided. Either way, *the use of the on-line information serves as feedback for building automation/maintenance so that this sort of anomaly can be avoided or handled more promptly.*

Building Temperature Distribution. In order to understand the spatial and temporal temperature distribution of 95 monitored rooms, we refine Figure 8 in terms of room daytime temperature and variation. we first plot Probability Mass Functions (PMFs) for daily temperatures across the year for all the rooms shown in Figure 9. Observe that the temperature difference can be as large as 18 degrees. Recall that the goal of HVAC system is to maintain each room temperature at $72^{\circ}F$ with ± 2 degree error, which is represented by the area between two vertical lines in the figure. *The wide span of spatial temperature difference indicates that spatial thermal condition in the building is non-uniform*, which motivates the high resolution characterization of thermal comfort and ill-conditioning in following subsections. Second, in terms of per-room temperature variation, we draw in Figure 10 the Cumulative Distribution Function (CDF) for standard deviation of daytime temperature over the year for all rooms. It is shown that 95% observations are within the expected 2 degree temperature error, indicating that the temporal temperature variation in Dreese is low. A small number of exceptions are probably due to HVAC mode change anomaly or human activities in the room such as accidentally leaving windows open, using personal heaters and etc.

Building Comfort Level. Given the target comfort range of $[70^{\circ}F, 74^{\circ}F]$, Figure 11 discriminates the frequency of three types of rooms in terms of their daily average temperatures: hot ($> 74^{\circ}F$), comfortable ($[70^{\circ}F, 74^{\circ}F]$), and cold ($< 70^{\circ}F$) rooms. We observe that the ratio of the comfortable area to the whole area is *on average* 47% *of the building*, i.e., more than half of the space is not well conditioned. This ratio is even lower for the Green Building LEED Silver Certification policy, which mandates maintaining temperatures during heating and cooling periods at $70^{\circ}F$ and $76^{\circ}F$ with ± 2 degree error respectively when occupied [10]. *These measures provide strong evidence for*

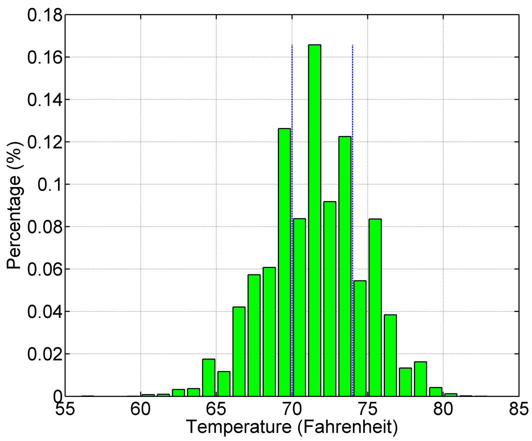

Figure 9. PMF of room daytime temperature over the year.

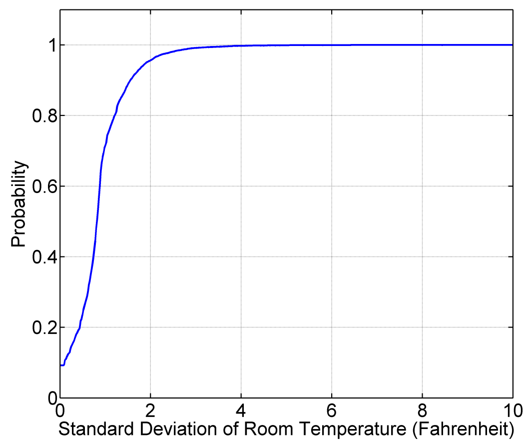

Figure 10. CDF for standard deviation of room daytime temperature over the year.

refining the global control policy of the building.

Ill-Conditioned Rooms. Analysis of individual room data reveals that there are some rooms for which adjusting the global control policy may not suffice to satisfy the thermal comfort objective. For instance, some rooms are always hot or cold no matter what the HVAC mode of operation is.

More specifically, the percentage of time that a room's daily temperature falls outside the acceptable range determines its level of ill-conditioning. To account for factors such as standard deviation of daily temperatures and accuracy errors in temperature sensor readings, we extend the comfort range by ± 2 degrees, i.e., if the daily average temperature of a room is beyond the range of $[68°F, 76°F]$, the room is considered to be ill-conditioned on that day.

Figure 12 plots the Cumulative Distribution Function (CDF) for percentage of days over a year that a room has been ill-conditioned. It is observed that approximately 80% of rooms are cumulatively over- or under-conditioned for less than 25% of the time, i.e., three months in

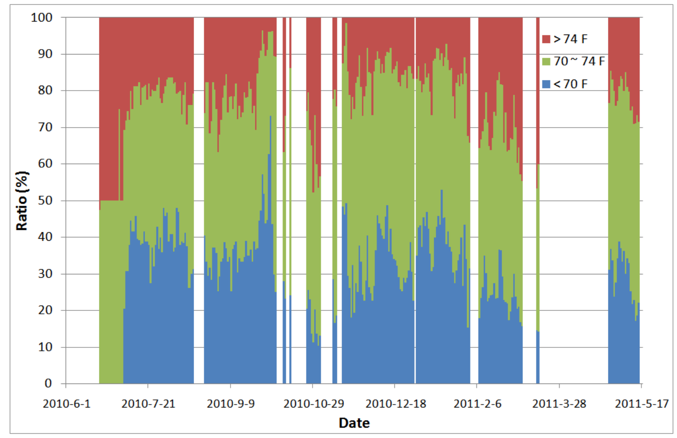

Figure 11. Daily ratios of hot ($> 74°F$), comfortable ($[70°F, 74°F]$), and cold ($< 70°F$) rooms in Dreese building.

a year; a few rooms are ill-conditioned for most of the year. Accordingly, Figure 12 classifies rooms into three categories: persistently ill-conditioned, intermittently ill-conditioned, and (marginally) well-conditioned rooms, whose ill-conditioning ratios are respectively in a range of $[60\%, 100\%]$, $[25\%, 60\%)$, and $[0, 25\%)$. Incidentally, we also observe that the average length of contiguous ill-conditioned periods is highly correlated with the cumulative number of ill-conditioned days.

In other words, a majority of rooms have temperatures which oscillate around the border of acceptable range from time to time. We consider them to well-conditioned or borderline rooms; the existence of borderline rooms is to be expected given the complexity of air distribution/diffusion in a large building; the number of these rooms can likely be controlled by proper selection of global control policies.

The existence of a number of intermittently ill-conditioned rooms is an emergent finding of our analysis. That most of these rooms remain ill-conditioned for non-trivial lengths of time suggests that these rooms may need parameter adjustment at the local controllers or even HVAC maintenance. As for persistently ill-conditioned rooms, maintenance of the HVAC system may be required.

Figure 13 visualizes the result of our classification in a 2-D map of Dreese Laboratory. The rows correspond to floor levels and rooms are displayed counter-clockwise with regard to the floor layout in Figure 3. Due to differences in floor plan and room size, some floors have fewer rooms than others. We distinguish persistently hot, persistently cold, intermittently hot, intermittently cold rooms by red, blue, pink, and cyan colors, respectively. Borderline and well-conditioned rooms are all colored in green since occasional thermal variations are difficult to avoid in the building. Figure 13 reveals that the number of cold rooms is larger than that of hot rooms.

We investigated a few ill-conditioned rooms and interviewed their occupants. Anecdotally, for instance, the staff

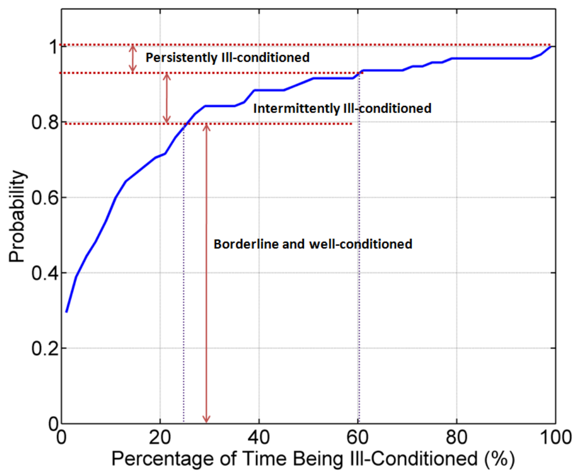

Figure 12. Room CDF for percentage of ill-conditioned days over the year.

Figure 13. Room classification in Dreese building —Red: persistently hot; Blue: persistently cold; Pink: intermittently hot; Cyan: intermittently cold; Green: borderline and well-conditioned rooms.

and students working in persistently hot rooms complained about feeling thermally uncomfortable year round. As for persistently cold rooms on the 8th floor, room 886 serves as the data center for CSE department, which is maintained separately at a rather low temperature, and room 883 is a small-sized office room at the corner of the floor. An interesting observation is that a nontrivial number of ill-conditioned rooms are located at corners of the building, such as room 779, 679, 883, 281, 791, 691, 591, 399, 898, and 798, which suggests that local controllers around corners need careful adjustment to ensure their effectiveness. Some of the persistently or intermittently cold rooms are relatively large-space inner rooms, such as 480, 395, and 380 (376 is a small room almost inside 380). Since almost every individual room's pneumatic controller can be adjusted locally (with a change in supply air temperature, supply static pressure, hot water temperature, and etc.), *the ill-conditioning classifier guides building operators to diagnose areas that demand local adjustment or maintenance, either remotely or physically.*

V. HVAC System Efficiency Assessment

This section analyzes opportunities for increasing the building's HVAC system efficiency at different time scales as well as with occupancy awareness.

Before discussing potential energy savings, we present energy profile for Dreese building by evaluating a representative building energy model developed using information from the building audit, operator interviews, and building plans. Chart in Figure 14 models Dreese Laboratory's energy consumption in eQuest simulator [20]. The electric energy is mostly consumed by fan and refrigeration cycle, equivalent to $18,407$ MBTU/year[3]. It is observed that steam use in winter for heating purpose is significantly higher than that in summer due to humidification, the total energy consumption being $16,292$ MBTU/year.

Reducing On-going Over-Conditioning. As seen in Figure 11, *about* 40% *of the building is over-conditioned (blue) in the cooling periods from July to October, and* 30% *of the building is over-conditioned (red) in the heating periods from November to May*. This substantial over-conditioning of the building suggests that the control policy for setting supply air discharge temperature can be made less aggressive. An eQuest analysis of the building shows that 17% energy can be saved by adopting appropriate supply air temperature reset[4].

Fine-tuning the set point can be performed experimentally since ThermoNet on-line information feeds can be readily integrated with on-line control policy enforcement. In fact, the existing AHU set points are determined based on only three statically selected rooms for achieving control objective. Limited samples from the building tend to result in over/under-conditioning. Fine grain instrumentation from ThermoNet enables utilizing the average or an appropriate fraction of room thermostats to guide set points.

Longer Temperature Setbacks. Figure 15 shows the diurnal thermal dynamics. Specifically, we compute for each room its hourly temperature averaged over a season and show the hourly indoor temperature as the mean of all rooms' hourly values for each season. Consistent with previous observations, the temperatures in Figure 15 are actually warmer during the winter and colder during the summer due to over-conditioning. Furthermore, the hourly temperature exhibits two waves during a day with peaks occurring around 5am and 5pm, respectively. These waves delineate thermal dynamics in the building associated with the HVAC

[3]BTU stands for British Thermal Unit, which is a unit of energy consumed by or delivered to a building. A BTU is defined as the amount of energy required to increase the temperature of 1 pound of water by 1 degree Fahrenheit, at normal atmospheric pressure.

[4]Even more energy would be saved by properly adjusting set points to conform to the Green Build policy.

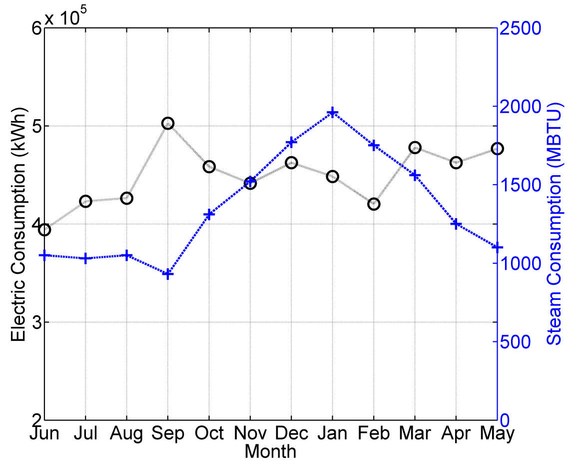

Figure 14. Modeled energy consumption in Dreese Laboratory.

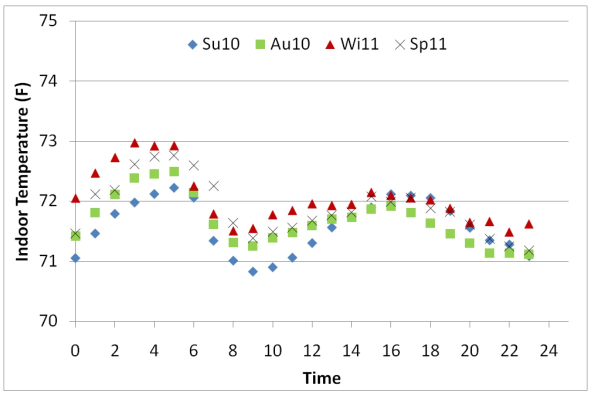

Figure 15. Indoor hourly temperatures of a day for four seasons.

cycling. Note that temperature variation between 12am and 7am is more substantial than that between 8am and 10pm as a result of AHU sleep during night-time, nevertheless, the overall average temperature variation is within only 2 degrees. During the day, the variation within 1 degree except for summer where the influence of outside temperature on indoor conditions becomes slightly significant. Note that temperature changes slowly at night when the AHUs are sleeping, and by 1-1.5 degrees within an hour of the AHUs resuming.

Figure 15 suggests that the AHUs can sleep more while not violating the building's control objective. Towards estimating the increase of sleep period, we characterize the occupancy of rooms for each hour by mining through illumination measurements acquired by ThermoNet. Basically, a room is assumed to be occupied if the light measurement is higher than a pre-determined threshold. ThermoNet sensors are installed carefully away from the window, which minimize the impact of outdoor illumination. However, the current occupancy estimation of Dreese does not distinguish the case that occupants do not turn off the lights when they leave the room. In the future work, adding motion detection sensors such as PIR and radar sensors to ThermoNet would

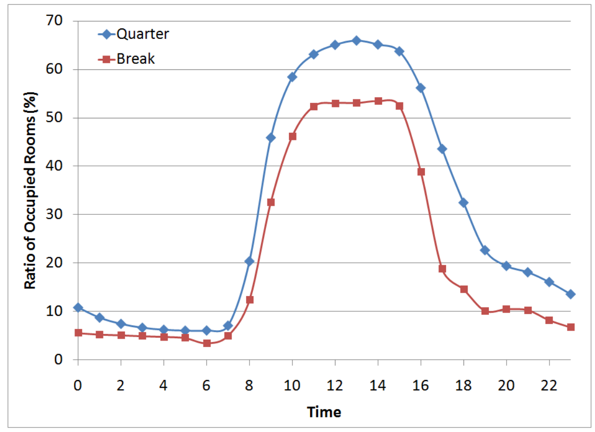

Figure 16. Percentages of occupied rooms for each hour of a day during one quarter and during the quarter break.

refine the estimation on room occupancy. Figure 16 presents the percentage of occupied rooms during one quarter and during the break preceding it, which verifies that the building is largely unoccupied from 10pm to 7am. Correlating this information with air diffusion rates, we *recommend setback periods to be* 10pm *to* 7am *over the quarter and* 8pm *to* 8am *over the break*. This schedule would respectively save 3 and 6 hours out of the 18 working hours every day during the quarter and break, which would correspond to 16.7% and 33.3% energy savings.

VI. Conclusions

Our long-term data-driven analysis based on ThermoNet reveals a significant scope for improving the thermal comfort of the building, in terms of refinement of its control policy as well as its realization, and for fine-grain local adjustment and maintenance associated with specific rooms. It also reveals a significant scope for improving energy efficiency. These suggest that there may be significant value associated with coupling the on-line information received from a long-lived WSN with the control system operation and facilities maintenance operation.

Issues. Although ThermoNet has successfully provided high-fidelity environmental information of the investigated building over a year, there are several methodological concerns that warrant further improvement. First, the precision of thermal monitoring can be increased by deploying multiple sensors in a room. Currently, only one sensor is installed per room no matter what size the room is, which raises the concern that the location of the sensor may be inappropriate since there may be air stratification problems that prevent it from making a proper measurement (since our sensors are typically co-located with the local room controller, we believe this is a risk only if the room controller was already located in a problematic location during building design). Also, a large-space room may need more than one sensor to characterize air dynamics at various locations, which has not been implemented in ThermoNet yet.

Second, multiple dimensions of physical information are necessary to better evaluate building thermal conditions. As discussed earlier, the thermal comfort standard PMV depends on many other physical information in addition to temperature. Deploying different types of sensors such as humidity and radiant temperature will enable us to estimate building PMV metric more accurately. On the other hand, the accuracy of occupancy prediction can be further improved by adding motion detection sensors, such as PIR or radar sensors, which could lead to a design of a more energy-efficient per-room control policy.

Third, a network level of sensor calibration deserves further exploration. Even though errors of temperature sensors were sample tested to be within 2 degrees, certain sensors may still have larger skews. Additionally, the accuracy of sensor outputs decreases as the circuit voltage supplied by batteries drops below certain threshold. In order to tolerate faults introduced by anomalous sensors, an in-field calibration scheme is in demand, which essentially discovers anomalies via observations from correlated sensors and adjusts the skew dynamically.

Improving HVAC system. ThermoNet reveals the fact that sampling from three statically selected rooms is insufficient for effectively achieving the control objective. One recommendation is to take into consideration of a larger sample of room thermal conditions. The scheme would likely reject samples from ill-conditioned rooms. An open question is how to sample from tens if not hundreds of rooms so as to maximize the effectiveness and efficiency of the HVAC system. We are currently in the process of discussing our findings and recommendations with building operators to design more effective and efficient HVAC strategies for the building being investigated.

Acknowledgment

Sincere thanks to: CSE Department staff members, Dave Kneisly, Mike Compton, Don Havard, and Aaron Jenkins, for helping us deploy ThermoNet in Dreese; OSU Building Automation Services colleagues, Kelly Bloomfield, Peter Calamari and Patrick Smith, and OSU Energy Services and Sustainability colleagues, Tracy Willcoxon, Gregory Roebke, and Aparna Dial, for sharing control system operation, objectives, and utilization data.

References

[1] V. L. Erickson and Cerpa, "OBSERVE: Occupancy-Based System for Efficient Reduction of HVAC Energy," ser. IPSN '11.

[2] G. Gao and K. Whitehouse, "The self-programming thermostat: optimizing setback schedules based on home occupancy patterns," ser. BuildSys '09, pp. 67–72.

[3] Y. Agarwal, B. Balaji, R. Gupta, J. Lyles, M. Wei, and T. Weng, "Occupancy-driven energy management for smart building automation," in *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, ser. BuildSys '10, pp. 1–6.

[4] V. L. Erickson and A. E. Cerpa, "Occupancy based demand response HVAC control strategy," ser. BuildSys '10, pp. 7–12.

[5] C. M. Liang, J. Liu, L. Luo, A. Terzis, and F. Zhao, "RACNet: a high-fidelity data center sensing network," ser. SenSys '09, pp. 15–28.

[6] A. Schoofs, A. G. Ruzzelli, and G. M. P. O'Hare, "Appliance activity monitoring using wireless sensors," in *Proceedings of the 9th ACM/IEEE Conference on Information Processing in Sensor Networks*, ser. IPSN '10, pp. 434–435.

[7] Z. C. Taysi, M. A. Guvensan, and T. Melodia, "TinyEARS: spying on house appliances with audio sensor nodes," ser. BuildSys '10, pp. 31–36.

[8] X. Jiang, M. Van Ly, J. Taneja, P. Dutta, and D. Culler, "Experiences with a high-fidelity wireless building energy auditing network," in *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '09, pp. 113–126.

[9] A. Rowe, M. Berges, and R. Rajkumar, "Contactless sensing of appliance state transitions through variations in electromagnetic fields," ser. BuildSys '10, pp. 19–24.

[10] "Green Build and Energy Policy, University Design Standards," ser. http://fod.osu.edu/ess.

[11] B. W. Olesen and K. C. Parsons, "Introduction to thermal comfort standards and to the proposed new version of en iso 7730," in *Energy & Buildings*, ser. July 2002.

[12] Aprilaire 8051 Flush Mount Temperature Sensor. http://www.smarthome.com/0657/Aprilaire-8051-Flush-Mount-Temperature-Sensor/p.aspx.

[13] TMP36 Voltage Output Temperature Sensors. http://www.analog.com/en/mems-sensors/digital-temperature-sensors/tmp36/products/product.html.

[14] THUM - USB Temperature/Humidity Sensor. http://www.smarthome.com/15250/THUM-USB-Temperature-Humidity-Sensor/p.aspx.

[15] Venstar Wireless Thermostat and Receiver. http://venstar.com/Thermostats/.

[16] Y. Agarwal, B. Balaji, R. Gupta, J. Lyles, M. Wei, and T. Weng, "Duty-cycling buildings aggressively: the next frontier in hvac control," ser. IPSN '11.

[17] STARGATE. https://www.eol.ucar.edu/rtf/facilities/isa/internal/CrossBow/DataSheets/stargate.pdf.

[18] TelosB. http://www.memsic.com/products/wireless-sensor-networks/wireless-modules.html.

[19] OARDC Weather System. http://www.oardc.ohio-state.edu/newweather/default.asp.

[20] eQuest simulator. http://doe2.com/equest/.