

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Shin, Eun-Jeong, Roberto Yus, Sharad Mehrotra, and Nalini Venkatasubramanian. "Exploring Fairness in Participatory Thermal Comfort Control in Smart Buildings." In Proceedings of the 4th ACM International Conference on Systems for Energy-Efficient Built Environments, 1–10. BuildSys '17. New York, NY, USA: Association for Computing Machinery, 2017. <https://doi.org/10.1145/3137133.3137156>.

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.



Exploring Fairness in Participatory Thermal Comfort Control in Smart Buildings

Eun-Jeong Shin

University of California, Irvine
ejshin1@uci.edu

Sharad Mehrotra

University of California, Irvine
sharad@ics.uci.edu

Roberto Yus

University of California, Irvine
ryuspeir@uci.edu

Nalini Venkatasubramanian

University of California, Irvine
nalini@ics.uci.edu

ABSTRACT

Current buildings rely on predefined rules to control the temperature in rooms disregarding their residents' thermal comfort. Multiple approaches have been presented in the literature to tackle this issue (e.g., by enabling occupants to express their feedback using their mobile devices). In general setting, to reach consensus among group members, we apply decision making methods (e.g., majority, mean, trimmed mean, and median) to aggregate the diverse expected comfort levels of the inhabitants. However, such methods might be *unfair* to some participants.

In this paper, we present the first study of the issue of fairness in participatory thermal comfort control. Inspired by the traditional definitions in scheduling, we introduce a definition of fairness that is suitably adapted to the particularities of our scenario. We then present our design of an aggregation method that ensures fairness. Finally, we show how our algorithm behaves compared with traditional aggregation methods in diverse simulated scenarios.

CCS CONCEPTS

• **Theory of computation** → **Scheduling algorithms**; • **Computing methodologies** → **Simulation tools**;

KEYWORDS

Thermal comfort, Personalized control, Fairness

ACM Reference Format:

Eun-Jeong Shin, Roberto Yus, Sharad Mehrotra, and Nalini Venkatasubramanian. 2017. Exploring Fairness in Participatory Thermal Comfort Control in Smart Buildings. In *Proceedings of 4th ACM Conference on Systems for Energy-Efficient Built Environments (BuildSys'17)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3137133.3137156>

1 INTRODUCTION

The control of the temperature in most of our current buildings (e.g., offices) is still based on rules that try to match the comfort of average people. For instance, Federal Occupational Safety and health Administration (OSHA) regulations usually recommend that

the temperature in the different rooms be in the 20-25°C range¹ which is considered to be comfortable. The reality is that people have diverse comfort preferences depending upon different factors such as their gender, age, weight, and nationality [19, 34]. Such prespecified control tends not to be comfortable for many of the inhabitants of the building [22].

Smart buildings equipped with a range of plug-and-play IoT devices that provide sensing/actuation capabilities are becoming the norm; these technologies allow users to interact with spaces they live/work in on a fine grained basis and participate in the control and management of these spaces. Participatory engagement with building environments in a more integrated manner enables a new level of personalized services including customized capture and exchange of information. Such participation can be used in the context of thermal comfort control. In this context, Participatory Comfort Control (PCC) systems adapt the temperature of buildings to increase the comfort of their residents.

However, adapting the temperature to the desired comfort level of the inhabitants of the building presents several challenges. On the one hand, there is a need for mechanisms to allow people to express their feedback regarding the temperature in the room they are. On the other hand, the feedback of the users have to be implemented in the HVAC (Heating, ventilation and air conditioning) system of the building. Multiple approaches have been presented in the literature to address these challenges. For instance, systems have been proposed to obtain feedback from people automatically (e.g., by determining their thermal profiles using information about their gender, age, height and weight [3] or manually (e.g., by enabling them to vote using their smartphones) [12, 21, 23]. Also, different mechanisms have been presented to translate the preferences of the users into the appropriate HVAC settings [5, 6, 9].

Nevertheless, there is an intrinsic challenge in such PCC systems: people inside a same room usually have different thermal comforts that have to be somehow aggregated to decide the temperature for the room (unfortunately, having different temperature zones in the same room might not be always possible). Works in the literature dealing with this problem often leverage traditional decision making methods to aggregate the possibly diverse opinions of people. For example, group decision making processes, such as majority or mean, have been extensively studied to reach consensus [31]. However, these methods might not be *fair* with the participants. For example, think about a group of people that meet daily in the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
BuildSys '17, November 8–9, 2017, Delft, Netherlands
© 2017 Association for Computing Machinery.
ACM ISBN 978-1-4503-5544-5/17/11...\$15.00
<https://doi.org/10.1145/3137133.3137156>

¹https://www.osha.gov/pls/oshaweb/owadisp.show_document?p_table=INTERPRETATIONS&p_id=24602

same room (e.g., coworkers sharing a meeting room). If a majority of the people prefer cold temperature and a minority prefer the temperature to be hot, the previously mentioned PCC systems might determine that the temperature in the meeting room should be cold always. In the long term, the minority will be clearly uncomfortable.

The focus of this work is the study of fairness in the aggregation of user thermal comfort in PCC systems. In this paper, we present the first study, up to the authors' knowledge, on this issue. In particular, the contributions of this paper are as follows:

- We introduce a definition of fairness, inspired by the traditional definitions in scheduling, adapted to the particularities of our scenario. Under this definition we prove that systems that use traditional group decision making processes, such as majority, mean, trimmed mean, and median, to aggregate people thermal comfort are not fair.
- We present our design of an aggregation algorithm that ensures fairness with the participants. We prove mathematically that the proposed algorithm is fair regardless of the number of people, their thermal comfort preferences, and iterations.
- We present a tool to simulate scenarios to evaluate different aggregation algorithms. The tool enables us to define the thermal profile of the inhabitants of a building, the frequency of the feedback obtained from them, and the desired aggregation algorithm, and executes a simulation that computes its fairness and comfort.
- We show experimental results of our approach and the traditional aggregation group decision making methods used in the literature (i.e., majority, mean, trimmed mean, and median) using the simulation tool with two scenarios: a weekly group meeting and a building with hundreds of participants and 64 zones. For the latter, we drive the simulation by the real occupancy data captured in a building.

The experimental results show the fairness of our algorithm in contrast to the rest. They also show that there exists a trade off between fairness and comfort. In order to guarantee fairness in any possible situation, the overall discomfort of the participants with our algorithm is higher, or similar in some situations, to the discomfort that results from other techniques that target minimizing discomfort without considering fairness.

The rest of the paper is structured as follows. In Section 2, we present the related work. In Section 3, we introduce a motivating scenario of a real building with different rooms and people. In Section 4, we present our modeling of fairness and comfort in occupant-participatory thermal comfort approaches. In Section 5, we show our fair algorithm for the aggregation of user thermal preferences. Finally, the experimental evaluation, conclusions and future work are presented in Section 6 and Section 7, respectively.

2 RELATED WORK

There have been multiple contributions in the literature towards participatory comfort control in buildings. However, up to the authors' knowledge, we are presenting the first approach that takes into account fairness in this context. In the following, we review works in the area of participatory thermal comfort control. In particular, we describe approaches to obtain user thermal comfort

preferences and to implement such preferences into the HVAC system. Also, we review works considering fairness in other aspects of computer science.

2.1 Works on Participatory Comfort Control

Thermal comfort is defined by American Society of Heating, Refrigerating and Air-conditioning Engineers (ASHRAE) Standard 55-2013 as the condition of mind that expresses satisfaction with the thermal environment. The *Predicted Mean Vote (PMV)* scale is a common mechanism to measure how comfortable a person is with room temperature [8]. The PMV scale consists of seven points of thermal sensation varying from cold (-3) to hot (3), with other points on the scale corresponding to *cool*, *slightly cool*, *neutral*, *slightly warm* and *warm*. Predicted Percentage of Dissatisfied (PPD) predicts the percentage of occupants that will be dissatisfied with thermal conditions. PPD is a function of PMV, given that as PMV moves further from neutral (0), PPD increases.

There are two main issues to be addressed in order to provide participatory thermal comfort: 1) Capture and aggregate user thermal preferences, and 2) Implement appropriate thermal comfort into the HVAC system. Several approaches for PCC systems that take into account both issues have been presented in the literature [2, 16, 22, 32]. Industry is also paying attention to the development of PCC systems. For instance, Comfy² and Honeywell's Vector³ systems collect user thermal feedback which is communicated to the building administrator. However, all these systems use traditional decision making methods (i.e., majority, mean, trimmed mean, and median) to aggregate user thermal preferences which, as we will demonstrate in Section 4.2, are not fair to the participants.

Some works have focused on obtaining the user thermal comfort and could serve as input for our aggregation algorithm. In [23], the authors developed a Temperature Comfort Correlation (TCC) model that builds a profile for each occupant based on three parameters: 1) personal information such as age, gender, height, weight, etc, 2) indoor and outdoor temperatures, 3) trainable parameters such as users votes. Thermal comfort is a function of heat gain and heat loss, which are primarily settled by metabolic rates and parameters are plugged into this function. In [29], the authors introduced a framework for building occupants serving as participatory sensors wherein occupants provide their feedback of thermal comfort and dynamic temperature control is applied [2, 6, 17, 30, 32] based on the feedback. With increasing levels of instrumentation in buildings, e.g., audiovisual sensing, methods to extract thermal comfort of occupants using visual signals and features along with specialized image processing have been studied [18].

The aggregated thermal comfort of the different users in a zone computed by our algorithm can be translated into HVAC control. Different strategies have been proposed in the literature for such a task. In [6], authors proposed a real-time and learned strategy for temperature control. [5] proposed an HVAC control strategy based on occupancy prediction and real time occupancy monitoring via a sensor network of cameras. In [9], signals measured directly on the body were used to infer user comfort and control

²<https://www.comfyapp.com/the-app/>

³<https://buildingsolutions.honeywell.com/en-US/solutions/Vector/Pages/default.aspx>

the air-conditioning system to direct air flow where it was needed. Techniques to estimate occupancy in buildings via sensors in conventional building management systems have been shown to be effective for zone-based HVAC scheduling and managing energy usage effectively while ensuring thermal comfort [1].

2.2 Fairness in Scheduling

The notion of fairness in scheduling has been addressed along multiple dimensions and has addressed a range of trade-offs - including fair share scheduling in OS, delay scheduling for clustered systems [33, 35], wireless networks [4, 13], shared memory architectures [14, 15, 25], and cloud computing [11, 27]. We focus on the treatment of fairness in scheduling algorithms, in particular in the context of a carpooling problem, which is related to our scenario.

The carpooling problem supposes that N people have decided to form a carpool to commute to office and a decision has to be made regarding who has to drive on a given day. Being selected as a driver is a burden and thus, the driver role has to be shared by the carpool members. In this context, to be fair, each person should be driving approximately $1/k$ of the time that she rides with $k - 1$ others. For example, if the carpool consists of A , B , and C , then A might be expected to drive $1/2$ of the times that she rides with only B or only C and $1/3$ of the times that she rides with both B and C .

Fagin and Williams [7] generalized this definition. Assume that at time t , A has participated in the carpool on b_2 days when exactly 2 persons participated in the carpool, on b_3 when exactly 3 persons participated, and so on. They defined A 's *ideal number of drives* as the number $(1/2)b_2 + (1/3)b_3 + \dots + (1/N)b_N$. Then, their formal definition of fairness considers a carpool scheduling algorithm as *fair* if "for each N (where N is the number of member of the carpool), there is a number P such that whatever the schedule of arrivals, it is the case that at each time t and for each carpool member A , the number of times that A has actually driven differs from his ideal number of drives in absolute value by no more than P ".

In Fagin and Williams scheduling algorithm, U is considered to be a value that represents the total cost of a trip. U is a least common multiple of $1, 2, \dots, m$ where m is the largest number of people who ever ride together at a time in the carpool. The ledger maintained by the algorithm contain one row for the date and one column for each participant. It starts with initial record with all 0's. If there are k participants on the given day, and A is the driver, then A 's entry is increased by $U(k - 1)/k$ units. For the rest of the users who do not drive, entry is decreased by U/k units. To decide the next driver, the algorithm chooses the person with the lowest entry in the table. [7] shows the above algorithm maintains fairness with a bounded deviation from the ideal number of drives.

3 MOTIVATING SCENARIO

Consider a future commercial office building where the HVAC (Heating, ventilation and air conditioning) system has been instrumented to capture resident feedback on thermal comfort; furthermore, let us assume that the system is capable of incorporating this input when adjusting temperatures for different rooms in the building. Note that any existing system can be used as a starting point for this instrumentation [6, 12, 20, 21, 26, 30]. Let us also suppose that the distribution of people at a given time in the rooms

of the our scenario is as shown in Figure 1(a) and that their thermal comfort at a given time is given in Table 1(b).

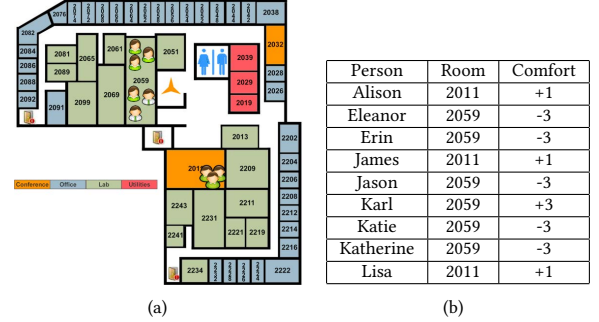


Figure 1: Distribution of the building residents in our sample scenario (a) and their thermal comfort (b).

In the case of rooms occupied by a single person, the individual's preferred temperature setting can be applied. When multiple residents of a space have identical thermal preferences (e.g., room 2011), the selection of a comfortable setting for all is also simple. When several residents in a room have varying thermal preferences (e.g., room 2059 where 5 users are cold (-3) and 1 of them feels hot (3)), one can apply any commonly used aggregation method to represent a group decision, such as: 1) *majority* (which selects the comfort supported by more than half of the participants), 2) *mean* (which selects the average value among votes), 3) *trimmed mean* (which selects the average value after trimming a percentage of the largest and smallest values), or 4) *median* (which selects the value, in an ordered set of values, below and above which there are an equal number of values). As an example, if we were to use the system presented in [28], the mean method would be applied to situations with several participant thermal comforts in the same room. A setting of "cold (-3)" (the mean of the comfort values) will result in discomfort for the single person that feels hot (3) at all times - such a smart building is *unfair* to some of its residents.

4 MODELING FAIRNESS AND COMFORT

In this section, we first explain the modeling of the participatory comfort control scenario. Then, we present a definition of fairness for such scenario that is inspired by the definition of fairness in the carpooling problem described earlier.

4.1 Modeling Thermal Comfort

Typically, PCC systems are slot-based control systems where decisions are made for a fixed interval of time, which we will refer in the following as a "round". We consider that at the beginning of each round, thermal comfort collection happens. We adopt the PMV scale (see Section 2.1) to model participant comfort. Thus, thermal comfort of a participant P_i at round r is a value from the PMV scale and can be defined as $TC_r[P_i] = [-3, 3]$ (see Table 1(b) for an example).

Let $P[z] = \{P_1, \dots, P_n\}$ be the set of participants in a specific zone z , and let $TC_r[P[z]]$ denote the aggregated comfort of zone

z set by the aggregate thermal control algorithm in round r . Note that $TC_r[P[z]] = [-3, 3]$ (i.e., it is also a value from the PMV scale). For example, considering the set of participants in room 2059 in our running example (see Figure 1(a)) and their thermal comfort (see Table 1(b)), if traditional aggregation methods such as majority, mean, trimmed mean, or median were used the value of $TC_r[P[2059]]$ would be -3, -2, -2, and -3, respectively.

Moreover, we can measure the disagreement between a participant's thermal comfort $TC_r[P_i]$ and the aggregated thermal comfort of the participants in the zone (viz., $TC_r[P[z]]$) as the difference in absolute value of these two terms $TCD_r[P_i] = |TC_r[P_i] - TC_r[P[z]]|$. In our running example, the disagreement per participant for the different approaches to aggregate thermal comfort discussed earlier is shown in Table 1. A measure of the total disagreement of the participants in the zone can be obtained by summing up their respective disagreements:

$$TCD_r[P[z]] = \sum_{i=0}^r TCD_r[P_i] \quad (1)$$

Table 1: Disagreement between the participant thermal comfort and the zone thermal comfort for traditional approaches in our running example.

Participant	Majority	Mean	Trimmed Mean	Median
Eleanor	0	1	0	0
Erin	0	1	0	0
Jason	0	1	0	0
Karl	6	5	6	6
Katie	0	1	0	0
Katherine	0	1	0	0

We can aggregate the disagreement of a user over time to obtain a value that represents how uncomfortable the user has been with the decisions made until round r as follows:

$$ATCD_r[P_i] = \sum_{k=0}^r TCD_k[P_i] \quad (2)$$

Finally, we generalize the scenario to buildings with m thermal zones where the temperature can be controlled. We define the set of zones as $Z = \{Z_1, \dots, Z_m\}$. Also, we define the set of participants of the building at a round r as $P = \{P_1, \dots, P_n\}$. P includes all the participants that have been in the building until round r .

4.2 Modeling Fairness

The definition of fairness in our context is based on the definition of fairness in the carpool scenario explained in Section 2.2. In the carpool scenario every time that users are involved in driving, they split the “cost” of the ride. A driver should ideally have driven at any point of time her share in the rides she took part in.

In our scenario, the equivalent of a ride is a round, which is the moment when the thermal comfort of the zone(s) is calculated using the thermal comfort of the N inhabitants within. During round r , each user P_i “loses” the value represented by her thermal comfort disagreement, which we refer to as *real loss* $L_r[P_i] = TCD_r[P_i]$. We define the “total loss” (or cost) of the round as $L_r = \sum_{i=0}^N L_r[P_i]$,

which is equal to the thermal comfort disagreement of the zone $TCD_r[P[z]]$ in Equation 1. Then, we split the loss L_r among the N users to define their *ideal loss*. We define the *extra loss* of a user P_i in round r , $EL_r[P_i] = L_r[P_i] - L_r/N$ as the difference between her ideal loss and her real loss. Finally, we define the *accumulated extra loss* of a user P_i at round r as:

$$AEL_r[P_i] = \sum_{k=0}^r EL_k[P_i] \quad (3)$$

We define that an algorithm to select the thermal comfort of a zone is “fair” if for the N users in the zone and at the round r , their accumulated extra loss in absolute value is bounded by some constant M which is independent of r .

Our definition (and subsequently algorithm) of fairness is based on the ones of the carpool scenario but is different. The main difference between the carpooling and the thermal comfort scenarios is that in the latter only one person is affected (the chosen driver) at the time of making a decision, whereas the rest of the people are benefited always and in the same way. In the thermal comfort scenario, users might have the same or similar preference, which would mean that selecting the temperature preferred by one user might also satisfy others. Therefore, straightforward algorithms (like round robin selection of the driver) or the greedy algorithm presented by Fagin and Williams [7] would not satisfy fairness in our context.

Under our definition of fairness it is possible to show that the traditional methods to aggregate participant thermal comfort are not fair.

Majority selection is not fair. Assume the scenario for room 2059 in Table 1(b) where there are six participants P_1, P_2, P_3, P_4, P_5 , and P_6 whose thermal comfort are -3, -3, -3, -3, -3, and 3, respectively. In this situation, selecting the value supported by the majority would make $TC_0[P[2059]] = -3$, the real loss per user will be 0, 0, 0, 0, 0, and 6 and the ideal loss per user will be 1. Thus, the extra loss will be -1, -1, -1, -1, -1, and 5. If this exact situation keeps happening for r rounds, the extra loss for participants P_1, P_2, P_3, P_4 , and P_5 will be $-1r$ and for participant P_6 will be $5r$, which are not bounded by any fixed number M (as r gets large).

Mean selection is not fair. Let's consider the same scenario described before. In that case, the mean value would make $TC_0[P[2059]] = -2$, the real loss will be 1, 1, 1, 1, 1, and 5, the ideal loss per user will be $5/3$, the extra loss will be $-2/3, -2/3, -2/3, -2/3, -2/3$, and $10/3$. Therefore, for r iterations the extra loss is again not bounded.

Trimmed mean selection is not fair. We can find an scenario similar to the previous for which this method is not fair regardless of the percentage removed from the smallest and largest user thermal comforts. For instance, using 20% (the value used in our experiments in Section 6) and the same scenario used for majority, the value selected will make $TC_0[P[2059]] = -3$. Therefore, the extra loss is the same than in the majority example and not bounded for r rounds.

Median selection is not fair. As in the case of the trimmed mean, median will make $TC_0[P[2059]] = -3$ and thus the extra loss of the participants is not bounded for r rounds.

5 A FAIR AGGREGATION ALGORITHM

In this section, we present our algorithm to calculate the thermal comfort of a zone using the thermal comfort of the participants within. Then, we prove that our algorithm is fair according to the definition of fairness explained in the previous section.

As explained earlier, both the thermal comfort of the participants and the zones take values in the range $[-3, 3]$. The goal of our algorithm is to find a group thermal comfort value within the range which will minimize the accumulated loss of the participants defined in Equation 3. To do that, our algorithm checks which participant had the highest accumulated loss in the previous round and decreases it by selecting a group thermal comfort value that will make his extra loss negative. So, if $AEL_r[P_i] > AEL_r[P_j] \forall j$ then $AEL_{r+1}[P_i] \leq AEL_{r+1}[P_j]$. Multiple group comfort values could be chosen to minimize the participants accumulated extra loss. Our algorithm tries to chose a value that would not increase the accumulated extra loss of other users with a different thermal comfort much. This is important as it can help in maintaining an appropriate bound on the accumulated extra loss in any situation. For that, the algorithm tries to find a group thermal comfort value that satisfies the two following conditions, whenever possible: 1) $AEL_{r+1}[P_i] \leq AEL_r[P_j] \forall j$; and 2) $AEL_{r+1}[P_i] > AEL_{r+1}[P_j] \forall j$. As we will show later, this can be achieved in situations where $TC_{r+1}[P_i] \leq TC_{r+1}[P_j] \forall j$.

Also, there might be situations where several participants share the same value for the accumulated extra loss and this value is the highest in a given round. For example, this happens the first iteration where $AEL_0[P_i] = 0 \forall i$. In this scenario, and from the point of view of fairness, any of such participants can be selected to decrease his accumulated extra loss. Our algorithm favors consistently the participant with the coldest thermal comfort preference as a way of breaking the tie.

In the following, we prove that our algorithm is fair for any scenario, number of participants and rounds, under the fairness definition given in Section 4.2.

Without loss of generality and in order to simplify the explanation of the proof, we will consider a simplified PMV scale $[-1, 1]$ for the participant thermal comfort. Under such scale, Table 2 shows the extra loss of the participants under the definition in Section 4.2 depending on their thermal comfort (i.e., $-1, 0, 1$), the number of participants supporting each thermal comfort (i.e., n_{-1}, n_0, n_1 with $n_{-1} + n_0 + n_1 = N$), and the aggregated thermal comfort computed (i.e., $TC_r[P[z]]$ or t in the table).

Table 2: Extra loss of a participant depending on his thermal comfort, number of participants supporting each thermal comfort (i.e., n_{-1}, n_0, n_1), and the aggregated thermal comfort computed (t).

	$TC = -1$	$TC = 0$	$TC = 1$
$t < 0$	$\frac{(2t+1)n_0 + 2tn_1}{N}$	$\frac{-n_1 - (2t+1)n_{-1}}{N}$	$\frac{n_0 - 2tn_{-1}}{N}$
$t > 0$	$\frac{n_0 + 2tn_1}{N}$	$\frac{-n_{-1} + (2t-1)n_1}{N}$	$\frac{-2tn_{-1} + (1-2t)n_0}{N}$

First, we prove three propositions that we will use to prove the fairness of our algorithm.

PROPOSITION 5.1. $L_r[P_i]$ is bounded by $[0, 2]$.

PROOF. $L_r[P_i]$ represents the loss for participant P_i when the thermal comfort of the zone selected is $TC_r[P[z]]$ that is, the difference between the participant's thermal comfort $TC_r[P_i]$ and $TC_r[P[z]]$. As both take values ranging between $[-1, 1]$, $L_r[P_i]$ is bounded by $[0, 2]$. \square

PROPOSITION 5.2. At any round r , $\sum_{i=1}^N EL_r[P_i] = 0$

PROOF. Using the extra loss model in Table 2 it is easy to see that multiplying each cell by its corresponding n_x (where x is $-1, 0, 1$), the sum of both rows is equal to zero. \square

PROPOSITION 5.3. If P_i and P_j are such that $AEL_r[P_i] = AEL_r[P_j]$ and $AEL_r[P_i] > AEL_r[P_k] \forall k$ and $TC_{r+1}[P_i] < TC_{r+1}[P_j]$ then $\exists TC_{r+1}[P[z]]$ such that $AEL_{r+1}[P_i] < AEL_r[P_i]$

PROOF. As $AEL_{r+1}[P_i] = AEL_r[P_i] + EL_{r+1}[P_i]$, in order for Proposition 5.3 to hold we should prove that we can find a value such that $EL_{r+1}[P_i] < 0$. According to the extra loss model in Table 2, if $TC_{r+1}[P_i] = -1$ this value will be such that $TC_{r+1}[P[z]] < -n_0/2(n_0 + n_1)$ and if $TC_{r+1}[P_i] = 0$ this value will be such that $TC_{r+1}[P[z]] < -(n_1 - n_{-1})/2n_{-1}$. \square

PROPOSITION 5.4. If $AEL_r[P_i] > AEL_r[P_j] \forall j$ and $TC_{r+1}[P_i] \leq TC_{r+1}[P_j] \forall j$, then $\exists TC_{r+1}[P[z]]$ such that $AEL_{r+1}[P_i] < AEL_r[P_i]$ and $AEL_{r+1}[P_i] > AEL_{r+1}[P_j] \forall j$.

PROOF. There exists two scenarios for $r+1$: $TC_{r+1}[P_i] = -1$ or $TC_{r+1}[P_i] = 0$ and $TC_{r+1}[P_j] \geq 0 \forall j$. Then, for 5.4 to hold $AEL_r[P_i] + EL_{r+1}[P_i] > AEL_r[P_j] + EL_{r+1}[P_j]$. In the first scenario and using the extra loss model of Table 2, $TC_{r+1}[P[z]] < -(AEL_r[P_i] - AEL_r[P_1])/2$ (with $AEL_r[P_1]$ representing the highest accumulated extra loss of a participant P_j whose $TC_{r+1}[P_j] = 1$). In the second scenario, $TC_{r+1}[P[z]] < -(AEL_r[P_i] - AEL_r[P_1])$. \square

Now we prove that our algorithm is fair.

THEOREM 5.5. $AEL_r[P_i] < 2N \forall i$

PROOF. We prove it by contradiction. We assume that at round r it was the first time that a participant P_i crossed the bound. Therefore,

$$AEL_r[P_i] > 2N \quad (4)$$

Because it was the first time, it means that in this round his extra loss made him increase his accumulated extra loss. By 5.1, we know that at max his extra loss could had been 2 so,

$$2N > AEL_{r-1}[P_i] > 2(N-1) \quad (5)$$

Also, our algorithm guarantees by proposition 5.3 and 5.4 that such increase could only happen in two situations:

- (1) $\exists P_j$ such that $AEL_{r-1}[P_j] > AEL_{r-1}[P_i]$ and $TC_r[P_j] > TC_r[P_i]$
- (2) $AEL_{r-1}[P_i] = AEL_{r-1}[P_j]$ and $TC_r[P_i] > TC_r[P_j]$

We can focus on the first situation, where P_j had a highest AEL than P_i and a hotter thermal comfort preference. In such situation, P_j would had decreased his AEL while increasing the AEL of P_i for a number of rounds. We can consider r' as the first round where $AEL_{r'}[P_j]$ started decreasing and

$$\begin{aligned} 2N &> AEL_{r'}[P_j] > 2(N-1) \\ 2N &> AEL_{r'}[P_i] > 2(N-2) \end{aligned} \quad (6)$$

Let's consider now the second situation, where P_j and P_i had the same accumulated extra loss at round $r-1$ and P_i had a hotter thermal comfort preference than P_j . Such tie should had been reached following the same idea than earlier, but this time being P_i the one that would had decreased his accumulated extra loss whereas P_j would had increased his. Therefore, 6 would interchange the P_i and P_j values.

In both situations and as explained before when discussing the increase on the accumulated extra loss of P_i , this could only happen if there was another participant P_k with an equal or higher accumulated extra loss at round $r'-1$ and a hotter thermal preference. In such scenario, we can also consider r'' as the first round when $AEL_{r''}[P_k]$ started decreasing and

$$\begin{aligned} 2N &> AEL_{r''}[P_k] > 2(N-1) \\ 2N &> AEL_{r''}[P_j] > 2(N-2) \\ 2N &> AEL_{r''}[P_i] > 2(N-3) \end{aligned} \quad (7)$$

We can keep using the same logic for all the possible participants. If we apply it recursively we will reach a time r'^n in which the last participant P_n would appear to make

$$\begin{aligned} 2N &> AEL_{r'^n}[P_n] > 2(N-1) \\ &\dots \\ 2N &> AEL_{r'^n}[P_j] > 2(N-(N-1)) \\ 2N &> AEL_{r'^n}[P_i] > 2(N-N) \end{aligned} \quad (8)$$

Which means that at that point all of them will have an accumulated extra loss greater than zero which cannot happen as it would contradict the condition on the first round where every participant satisfies $AEL_0[P_x] = 0 \forall x$. \square

6 EXPERIMENTAL EVALUATION

In this section, we first present the simulator tool developed to perform the experiments. Then, we explain the experimental set up with the scenarios. Finally, we show the results of the experiments to compare traditional aggregating methods and our approach in terms of fairness and comfort.

6.1 Simulator Tool

We have developed a simulator to test various building occupants scenarios. The simulator takes input parameters such as a set of zones, a set of participants with their corresponding thermal profiles (e.g. hot/neutral/cold preferred) and movement between rooms, a set of initial temperatures per room, simulation period, etc. Also, it enables the user to select the algorithms to use to compute thermal comfort of the rooms (it supports the five algorithms considered along the paper –majority, mean, trimmed mean, median, and our

algorithm–). Also, we included a variant of our algorithm which bases the selection of the participant whose thermal comfort should be selected at a given round on the L_{max} distance of the accumulated extra losses of the participants. We included this algorithm because it presents an interesting trade off fairness vs. comfort. Nevertheless, such algorithm is not fair according to our definition as we have found a pathological situation where the accumulated extra loss is unbounded⁴.

Once the settings have been defined, the tool starts the simulation of the movement of people through the different rooms and periodically uses the defined algorithm to compute the thermal comfort of the zone/room based on the thermal comfort of the participants in it. Finally, the tool simulates a change of temperature in such rooms according to the computed thermal comfort.

The simulator is composed of four main modules:

- Participant manager, which generates participants and their thermal comfort based on the profiles and ratio given as input (e.g., if the input ratio is 50:30:20, 50% of cold, 30% of neutral, and 20% of hot preferred participants will be generated). Also, it manages the movement of the participants according to the input trajectories.
- Zone thermal comfort manager, which computes the thermal comfort for each zone using the thermal comfort of the participants in the room and the selected algorithms.
- Fairness and thermal comfort calculator, which computes the accumulated loss and discomfort of each participant at each round.
- Temperature manager, which simulates a change of temperature according to the computed thermal comfort for each zone. HVAC control is out of the scope of the paper but we implemented a simple mechanism in the simulator that changes the temperature by 1°C every 30, 45, and 60 minutes if the computed thermal comfort is -3/3, -2/2, and -1/1, respectively.

After each simulation, the tool outputs CSV files with the results for each algorithm in terms of accumulated loss and discomfort per user and per round.

6.2 Experimental Setup

We used the simulator to empirically test the behavior of our approach compared with traditional aggregating methods. We tested two scenarios:

- (1) A simulation of a weekly group meeting with regular participants. We assume that 10 people meet in the same room every week at 8am. Meeting lasts for about 30 minutes and the comfort of the room is computed at the beginning of the meeting. The scenario tries to simulate short meetings, which are common in buildings like the one described in our experiments. Due to the time needed to implement the temperature changes, we compute the comfort at the beginning in order to change the temperature for the duration of the whole meeting. Also, we introduce some randomness for each meeting guaranteeing that at least 4 of the participants vote in each round.

⁴Due to the space restrictions we included more information about such situation and the complete algorithm at tippersweb.ics.uci.edu/fairness.

- (2) A complete building with hundreds of participants. The TIPPERS [24] dataset consists of a trace of distinct devices connected to 64 Wi-Fi access points located in the 6 story Bren Hall building at UC Irvine collected over a 20 month period from January 2016 to September 2017. We used a subset of the TIPPERS dataset for the period of May 29th to June 2nd 2017 from 8am to 8pm. To filter out passerby devices (some of the Wi-Fi AP ranges cover a few meters outside of the building), we discarded MAC addresses that had less than 50 events registered in the log for the period. After the filter, we obtained 18,837 events and 1065 unique MAC addresses. Wi-Fi AP ranges cover several rooms in most of the cases so, for simplicity, we considered each Wi-Fi AP as a zone. Therefore, all the MAC addresses connected to a Wi-Fi AP at a given time represent participants in the same zone.

We defined three types of thermal profiles for the two scenarios: hot preferred, neutral, and cold preferred (see Figure 2). These profiles are based on the temperature in the zone. For example, as Figure 2 shows, for a room temperature of 23°C the thermal comfort of a hot preferred, neutral, and cold preferred users will be -1 (slightly cold in the PMV scale), 0 (neutral), and 2 (warm), respectively. For each scenario, we performed two tests with two different thermal profile settings. One setting represents a scenario with a clear majority of thermal profiles, where 90% of people prefer cold temperature and 10% of people prefer hot temperature (in the following we refer to it as 90/10). The second setting represents a balanced scenario where 50% of people prefer cold temperature and 50% of people prefer hot temperature (50/50).




Figure 2: Thermal comfort for different thermal profiles in our experiments.

6.3 Results for the Meeting Simulation

In the following, we show the results of fairness and comfort achieved by each method for the first scenario.

6.3.1 Fairness.

Figure 3 shows a comparison of the fairness of the different algorithms in terms of accumulated loss per round. For each algorithm we show two graphs (for the 90/10 and 50/50 settings) and a measure of the average accumulated loss per participant (represented by the dots in the graph) and the standard deviation from it. Notice that for each graph, the value at the last round represents the final accumulated loss at the end of the simulation.

The first implication of the results is that the accumulated loss for the majority (Figure 3(a)), mean (Figure 3(b)), trimmed mean (Figure 3(c)), and median (Figure 3(d)) are not bounded. The different graphs show that the accumulated loss keeps increasing for each round for both settings. It is interesting to notice that the behavior is similar for both settings even when one of them represents a clear

majority vs minority situation. This is caused by the randomness introduced in terms of the participants at the meeting at each round. This means that even in the 50/50 setting, there might be rounds with a clear majority. If we focus on the results for our algorithm (see Figure 3(e)), we can see that, in contrast with the other algorithms, it keeps the accumulated loss bounded as expected.

Additionally, Figure 3(f) shows the variant of our algorithm based on L_{max} distance which also keeps the loss bounded in this test (although it is not bounded in general as commented before).

6.3.2 Thermal Comfort.

Figure 4 shows the comparison of the algorithms in terms of comfort (measured as aggregated thermal discomfort of the users per round). Figure 4(a) shows the comparison in terms of average discomfort per user for the setting 50/50 and we can see that as expected, the majority algorithm achieves the best comfort as it selects the most comfortable temperature for the majority of the participants. Our algorithm, increases the average discomfort in comparison with the traditional approaches. This happens because by being fair, sometimes it selects the thermal comfort of the minority and the majority of users thus, feel uncomfortable. In terms of the maximum value of discomfort (see Figure 4(b)), our algorithms achieve the worst comparing to the majority and median algorithms. For the 90/10 setting, our algorithm increases the discomfort greatly in comparison with others because every time the comfort of the minority gets selected, most of the participants feel uncomfortable. Interestingly, the variant of our algorithm using L_{max} behaves similarly to the rest of algorithms in terms of comfort for the 50/50 scenario.

6.4 Results for the Building Simulation

In the following, we show the results for the second scenario.

6.4.1 Fairness.

Due to the complexity of the scenario and the high number of participants in it (1065), we show test results achieved in terms of accumulated loss at the final round of the simulation. Table 3 shows accumulated loss for each method in both 90/10 and 50/50 settings. As in previous scenario, our algorithm achieves better results in terms of accumulated loss than most of the traditional approaches. Notice that for some specific cases, our algorithm obtains a slightly higher value than the mean/trimmed mean approach. We analyzed the data and discovered that the occupancy data extracted from the real building shows a situation where only in very specific moments the zones contain a significant amount of participants (e.g., more than 4) and in most of the rounds, most of the participants are in one or two zones (remember that the zones correspond to Wi-Fi APs). This corresponds to areas of high concentration of people in the building and to specific events.

6.4.2 Thermal Comfort.

In terms of comfort, we again looked at the accumulated discomfort at the final round of the simulation. Comparing the average discomfort per user, the results show that for the 50/50 setting, our algorithm achieves close discomfort compared to the majority algorithm which obtains the lowest discomfort (0.25 vs 0.2). In case of the 90/10 setting, the difference is 0.05, which is similar to the previous scenario (0.11 vs 0.06). Notice that as we pointed

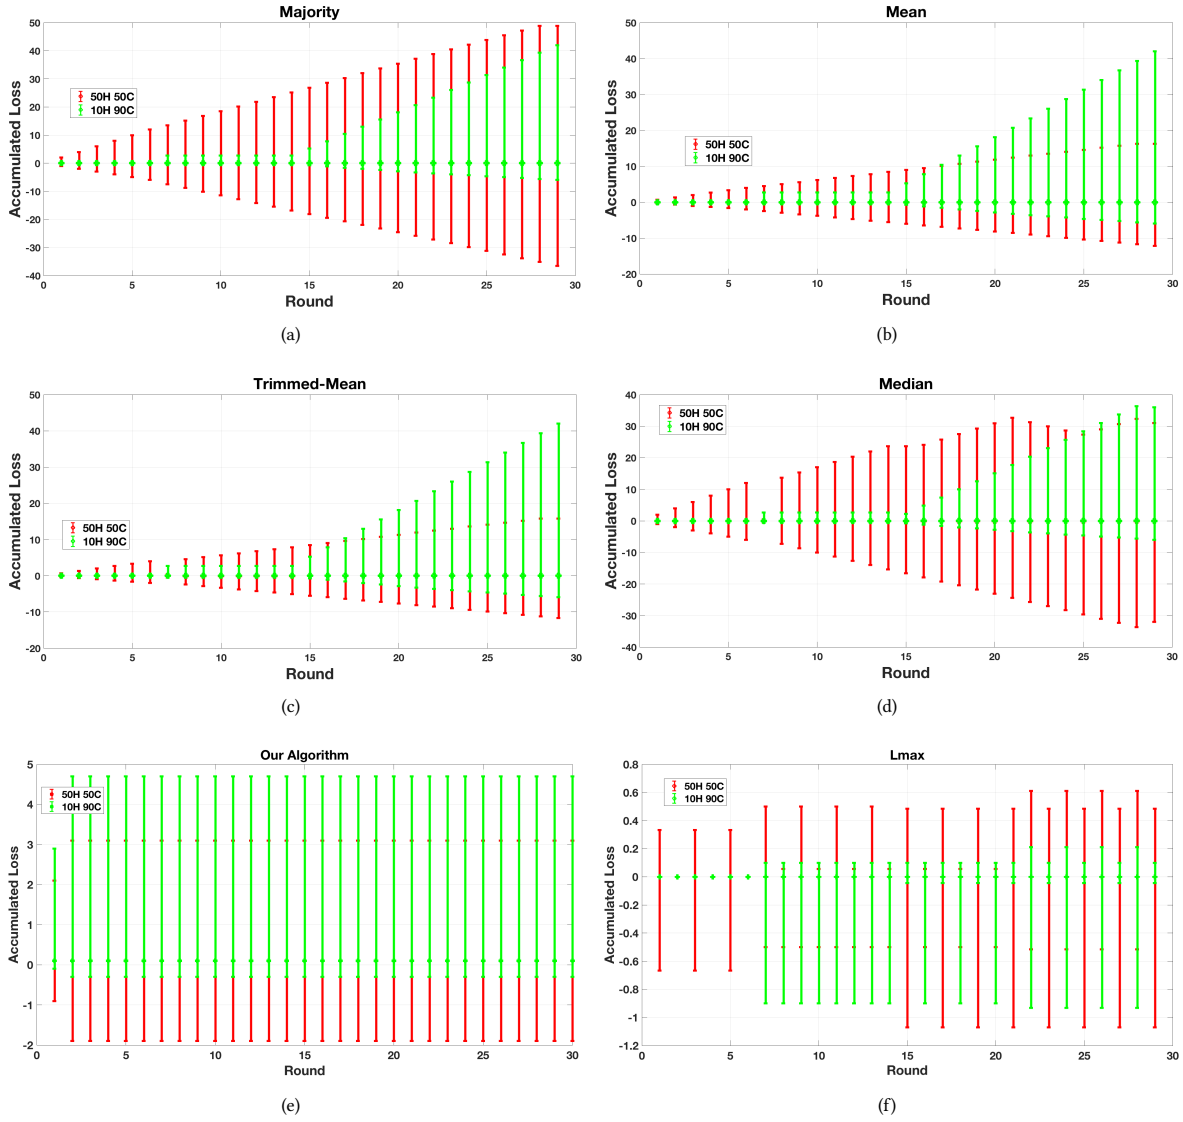


Figure 3: Comparison of fairness for the first scenario in terms of accumulated loss for traditional methods: majority (a), mean (b), trimmed mean (c), and median (d), our approach (e), and a variant (f).

Table 3: Accumulated loss at the last round for each algorithm in the building scenario.

Scenario	Loss	Majority	Mean	Median	Trim. Mean	Our Algo	Lmax Algo
90/10	Max	7.5	7.5	7.5	2.5	3	0.5
90/10	Min	-3	-1.5	-3	-1	-5.33	-0.83
50/50	Max	8.5	2.83	6	2.83	3	0.83
50/50	Min	-6.5	-2.16	6	-2.16	-5.33	-0.66

out in the previous scenario, the variants of our algorithm using *Lmax* keep the discomfort similar to the rest of algorithms. For the 50/50 setting, *Lmax* obtain an average discomfort of 0.48 and for the 90/10 setting, they obtain an average discomfort 0.8.

In summary, the experiments of the two scenarios and the different settings show that, as expected, our algorithm obtains the lowest accumulated loss per user while increasing the accumulated discomfort. An interesting conclusion of the experiments is that

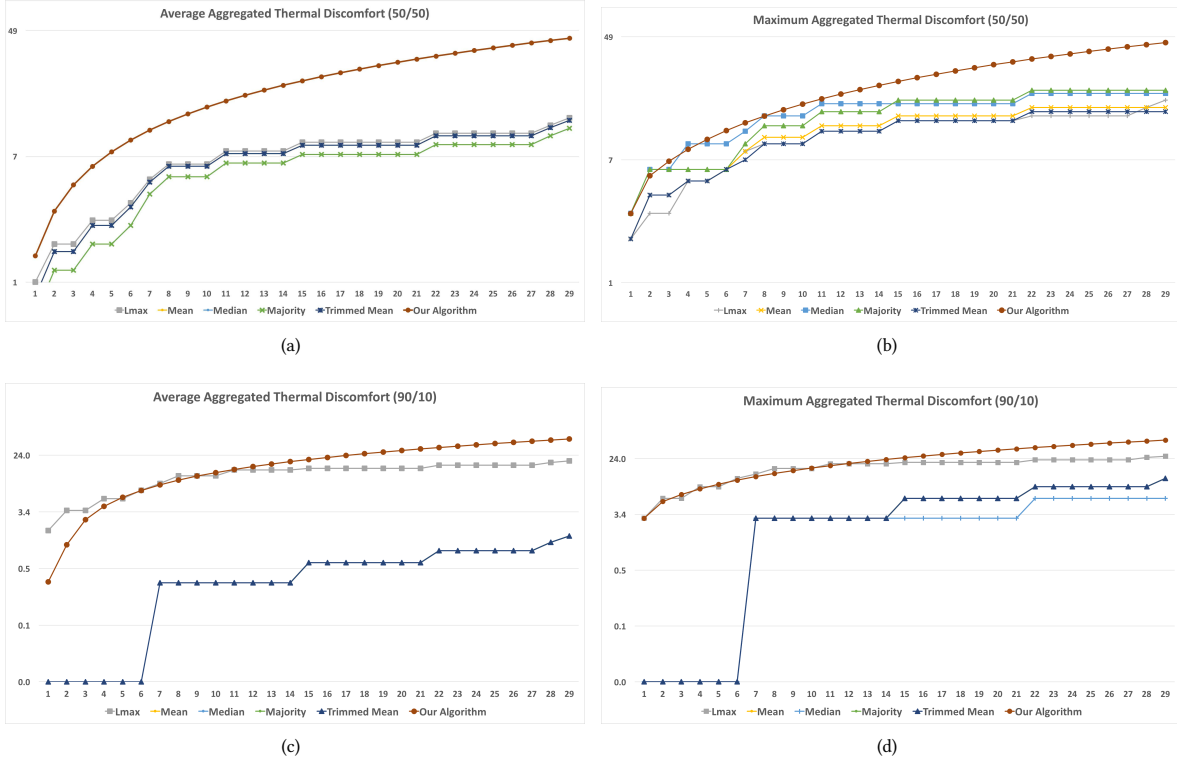


Figure 4: Comparison of comfort for the first scenario in terms of aggregated thermal comfort.

the L_{max} variant of our algorithm (which is not fair under our definition) maintains a low accumulated loss while maintaining the discomfort in similar levels than the traditional algorithms. Therefore, it seems that this variant could present a good trade off between fairness and comfort. However, it is not straightforward to design an algorithm that can be proven fair for any scenario while including the goal of maximizing the comfort of the participants.

7 CONCLUSIONS AND FUTURE WORK

Current buildings and their residents are equipped with the technology to make it possible to control the temperature in their rooms to maximize resident's comfort. Different proposals have been presented in the literature to address this challenge. In situations where more than one resident is in the room, there is a need for mechanisms to determine the temperature for the room according to their possibly diverse thermal comforts (e.g., some users prefer colder or hotter temperatures). The systems presented in the literature use traditional decision making methods (i.e., majority, mean, trimmed mean, median, etc.) to make such decision. We argue that these methods are not fair with the residents.

We have presented the first analysis, up to the authors' knowledge, of the fairness of the different methods used to aggregate residents' thermal comforts. Specifically, we defined fairness in the context of thermal comfort control and proved that traditional group decision making processes, such as majority, mean, trimmed mean, and median, to aggregate people thermal comfort are not fair.

We also explained an aggregation algorithm that ensures fairness with the participants. Our algorithm takes the thermal comfort of the participants as input but is agnostic about how this information is determined; it could be obtained from personal information (such as age, gender, height, and weight) [23], from stochastic thermal profiles related to room conditions [10], or even asked directly to the participants [29].

We presented experimental results performed using a simulator that we implemented to recreate scenarios driven by real occupancy data. The experimental results show how the accumulated loss for participants in our simulated scenarios increased with traditional decision making algorithms whereas it remains bounded for our algorithm. Also, they show, even when there exists a trade off between fairness and comfort, different variants of our algorithm can achieve a similar comfort than the traditional approaches.

As future work, we plan to continue our study about the trade off between fairness and comfort using other variants of our algorithm. The challenge is to prove that these variants, which achieved good results in terms of fairness and comfort in our experiments, are fair for any scenario. Another aspect which was out of the scope of the paper but in our plan is to study whether these algorithms can be gamed by participants. Achieving energy efficiency while obtaining thermal comfort has been addressed in many works. We also plan to study how to combine energy efficiency into our fairness vs. comfort trade off. Finally, it would be interesting to study how to generalize the concept of fairness for other participatory services

in the smart buildings (e.g., time professors spent with students in tutoring hours).

ACKNOWLEDGMENTS

This material is based on research sponsored by DARPA under agreement number FA8750-16-2-0021. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

Also, we would like to thank Dhrubajyoti Ghosh for his help with the execution of the experiments.

REFERENCES

- [1] Omid Ardakanian, Arka Bhattacharya, and David Culler. 2016. Non-Intrusive Techniques for Establishing Occupancy Related Energy Savings in Commercial Buildings. In *3rd ACM International Conference on Systems for Energy-Efficient Built Environments (BuildSys'16)*. 21–30.
- [2] Bharathan Balaji, Hidetoshi Teraoka, Rajesh Gupta, and Yuvraj Agarwal. 2013. ZonePAC: Zonal Power Estimation and Control via HVAC Metering and Occupant Feedback. In *5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings (BuildSys'13)*. 18:1–18:8.
- [3] Nuala M Byrne, Andrew P Hills, Gary R Hunter, Roland L Weinsier, and Yves Schutz. 2005. Metabolic equivalent: one size does not fit all. *Journal of Applied physiology* 99, 3 (2005), 1112–1119.
- [4] Prasanna Chaporkar, Koushik Kar, Xiang Luo, and Saswati Sarkar. 2008. Throughput and fairness guarantees through maximal scheduling in wireless networks. *IEEE Transactions on Information Theory* 54, 2 (2008), 572–594.
- [5] Varick L Erickson and Alberto E Cerpa. 2010. Occupancy based demand response HVAC control strategy. In *2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building (BuildSys'10)*. 7–12.
- [6] Varick L. Erickson and Alberto E. Cerpa. 2012. Thermovote: Participatory Sensing for Efficient Building HVAC Conditioning. In *4th ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys'12)*. 9–16.
- [7] Ronald Fagin and John H. Williams. 1983. A Fair Carpool Scheduling Algorithm. *IBM Journal of Research and Development* 27, 2 (1983), 133–139.
- [8] Poul O Fanger. 1972. *Thermal comfort. Analysis and applications in environmental engineering*. McGraw-Hill.
- [9] Mark Feldmeier and Joseph A Paradiso. 2010. Personalized HVAC control system. In *2010 Internet of Things (IOT)*. 1–8.
- [10] Ali Ghahramani, Chao Tang, and Burcin Becerik-Gerber. 2015. An online learning approach for quantifying personalized thermal comfort via adaptive stochastic modeling. *Building and Environment* 92 (2015), 86–96.
- [11] Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. 2011. Dominant Resource Fairness: Fair Allocation of Multiple Resource Types. In *8th USENIX Symposium on Networked Systems Design and Implementation (NSDI'16)*. 24–24.
- [12] Lam Abraham Hang-yat and Dan Wang. 2013. Carrying my environment with me: A participatory-sensing approach to enhance thermal comfort. In *5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*. 1–8.
- [13] Xiao Long Huang and Brahim Bensaou. 2001. On max-min fairness and scheduling in wireless ad-hoc networks: analytical framework and implementation. In *2nd ACM International Symposium on Mobile Ad Hoc Networking & Computing (MobiHoc'01)*. 221–231.
- [14] Sungjin Im, Janardhan Kulkarni, and Benjamin Moseley. 2015. Temporal fairness of round robin: competitive analysis for Lk-norms of flow time. In *27th ACM symposium on Parallelism in Algorithms and Architectures (SPAA'15)*. 155–160.
- [15] Sungjin Im, Janardhan Kulkarni, and Kamesh Munagala. 2014. Competitive Algorithms from Competitive Equilibria: Non-clairvoyant Scheduling Under Polyhedral Constraints. In *46th Annual ACM Symposium on Theory of Computing (STOC'14)*. 313–322.
- [16] Farrokh Jazizadeh and Burcin Becerik-Gerber. 2012. Toward Adaptive Comfort Management in Office Buildings Using Participatory Sensing for End User Driven Control. In *4th ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys'12)*. 1–8.
- [17] Farrokh Jazizadeh and Burcin Becerik-Gerber. 2012. Toward adaptive comfort management in office buildings using participatory sensing for end user driven control. In *4th ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys'12)*. 1–8.
- [18] Farrokh Jazizadeh and S. Pradeep. 2016. Can Computers Visually Quantify Human Thermal Comfort?: Short Paper. In *3rd ACM International Conference on Systems for Energy-Efficient Built Environments (BuildSys'16)*. 95–98.
- [19] Boris Kingma and Wouter van Marken Lichtenbelt. 2015. Energy consumption in buildings and female thermal demand. *Nature Climate Change* 5, 12 (2015), 1054–1056.
- [20] K Konis and M Annavaram. 2017. The occupant mobile gateway: a participatory sensing and machine-learning approach for occupant-aware energy management. *Building and Environment* 118 (2017), 1–13.
- [21] Andrew Krioukov and David Culler. 2012. Personal building controls. In *11th International Conference on Information Processing in Sensor Networks (IPSN'12)*. 157–158.
- [22] Abraham Hang-yat Lam, Yi Yuan, and Dan Wang. 2014. An Occupant-participatory Approach for Thermal Comfort Enhancement and Energy Conservation in Buildings. In *5th International Conference on Future Energy Systems (e-Energy'14)*. 133–143.
- [23] Abraham Hang-yat Lam, Yi Yuan, and Dan Wang. 2014. An occupant-participatory approach for thermal comfort enhancement and energy conservation in buildings. In *5th International Conference on Future energy systems (e-Energy'14)*. 133–143.
- [24] Sharad Mehrotra, Alfred Kobza, Nalini Venkatasubramanian, and Siva Raj Rajagopalan. 2016. TIPPERS: A privacy cognizant IoT environment. In *IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops)*. 1–6.
- [25] Onur Mutlu and Thomas Moscibroda. 2008. Parallelism-aware batch scheduling: Enhancing both performance and fairness of shared DRAM systems. *ACM SIGARCH Computer Architecture News* 36, 3 (2008), 63–74.
- [26] Marco Pritoni, Kiernan Salmon, Angela Sanguinetti, Joshua Morejohn, and Mark Modera. 2017. Occupant thermal feedback for improved efficiency in university buildings. *Energy and Buildings* 144 (2017), 241–250.
- [27] Qifan Pu, Haoyuan Li, Matei Zaharia, Ali Ghodsi, and Ion Stoica. 2016. FairRide: Near-Optimal, Fair Cache Sharing. In *8th USENIX Symposium on Networked Systems Design and Implementation (NSDI'16)*. 393–406.
- [28] Sean Purdon, Branislav Kusy, Raja Jurdak, and Geoffrey Challen. 2013. Model-free HVAC control using occupant feedback. In *IEEE 38th Conference on Local Computer Networks Workshops (LCN Workshops)*. 84–92.
- [29] Sasank Reddy, Deborah Estrin, and Mani Srivastava. 2010. Recruitment framework for participatory sensing data collections. In *8th International Conference on Pervasive Computing (Pervasive'10)*. 138–155.
- [30] Angela Sanguinetti, Marco Pritoni, Kiernan Salmon, Alan Meier, and Joshua Morejohn. 2017. Upscaling participatory thermal sensing: Lessons from an interdisciplinary case study at University of California for improving campus efficiency and comfort. *Energy Research & Social Science* (2017).
- [31] Elizabeth Taylor, Kelsey Hewitt, and WF Lawless. 2013. Group decision-making: Consensus rule versus majority rule. *Procedia Technology* 9 (2013), 498–504.
- [32] Daniel A Winkler, Alex Beltran, Niloufar P Esfahani, Paul P Maglio, and Alberto E Cerpa. 2016. FORCES: feedback and control for occupants to refine comfort and energy savings. In *2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp'16)*. 1188–1199.
- [33] Chee Siang Wong, Ian Tan, Rosalind Deena Kumari, and Fun Wey. 2008. Towards achieving fairness in the Linux scheduler. *ACM SIGOPS Operating Systems Review* 42, 5 (2008), 34–43.
- [34] N Yamtraipat, J Khedari, and J Hirunlabh. 2005. Thermal comfort standards for air conditioned buildings in hot and humid Thailand considering additional factors of acclimatization and education level. *Solar Energy* 78, 4 (2005), 504–517.
- [35] Matei Zaharia, Dhruba Borthakur, Joydeep Sen Sarma, Khaled Elmeleggy, Scott Shenker, and Ion Stoica. 2010. Delay scheduling: a simple technique for achieving locality and fairness in cluster scheduling. In *5th European Conference on Computer Systems*. 265–278.