smartSDH: An Experimental Study of Mechanism-Based Building Control

Ioannis Konstantakopoulos, Kristy Hamilton, Yashaswini Murthy, Tanya Veeravalli, Costas Spanos, *Fellow, IEEE*, and Roy Dong

***Abstract*—As Internet of Things technologies are increasingly being deployed, situations frequently arise where multiple stakeholders must reconcile preferences to control a shared resource. We perform a five-month long experiment dubbed "smartSDH" (carried out in 27 employees' office space) where users report their preferences for the brightness of overhead lighting. smartSDH implements a modified Vickrey–Clarke–Groves (VCG) mechanism. smartSDH assesses the feasibility of the VCG mechanism in the context of smart building control and evaluated smartSDH's effect using metrics such as light level satisfaction, incentive satisfaction, and energy consumption. After the initial eight weeks, we noted that average satisfaction with the light levels dropped from 3.44 to 2.94 on a 5-point Likert scale and that satisfaction with incentives dropped from 3.57 to 3.16. This indicates a participation burnout not typically captured by theoretical analyses of VCG mechanisms. Additionally, our experiment reduced energy consumption by 35.22% over a five-month period, without directly incentivizing lower energy consumption, which provides evidence that much of the current energy consumption may be wasteful. Finally, we observed that environmental readings had statistically significant relationships with the lighting preferences of users, indicating promise for learning preferences from other observable factors.**

***Index Terms*—Experimental validation, Internet of Things (IoT), mechanism design, smart building control, Vickrey–Clarke–Groves (VCG) mechanism.**

I. Introduction

RECENT advances in sensing, actuation, and communication technologies have allowed an unprecedented level of control over the behavior of our devices and an unprecedented fidelity of information about the state of our systems. This has found application in a wide variety of "smart" decision-making processes, including the management of parking spaces, the monitoring of water usage, and the energy-efficient operation of instrumented homes and offices [1].

However, as the set of application domains for these "smart" algorithms grow, we find ourselves in scenarios where multiple stakeholders care about the state of a shared resource. For example, consider the control of lights in an office space. At home, one can easily adjust the light settings in a nonproblematic way, as they are dictators of their own home's illumination. In contrast, for an office setting where multiple coworkers have different preferences about the intensity of lights, the definition of what "should" happen is not as clear. Put another way, as we start to look at applications with a shared resource, *we need to find ways to reconcile the opinions of multiple users.*

Building on our previous example, we can think of a mechanism that asks each person: "How bright do you want the lights?" Suppose there are two office occupants and our mechanism simply implements the average of the two votes. Furthermore, suppose your coworker voted for 100% intensity, and you wanted the lights at 75% intensity. In this situation, you would vote for 50% intensity, as the mechanism would implement the average, which would be your desired lighting level. This simple thought experiment reveals a crucial point: humans have an incentive to *strategically* report data to push mechanisms to their selfishly desired outcomes.

This problem is often studied in economics under the title of "mechanism design." The Vickrey–Clarke–Groves (VCG) mechanism [2], [3] is one of the most celebrated achievements in this area. For more information about VCG mechanisms, we refer the reader to the wonderfully accessible [4]. VCG mechanisms have many desirable properties. First, VCG mechanisms implement the most socially beneficial outcome. Second, all participants are incentivized to truthfully reveal their preferences. Third, the mechanism can be designed such that every participant is better off participating in the mechanism, either due to the outcome chosen or due to some endogenously determined payment.

The most common application of VCG mechanisms is in auction-like settings. In these settings, people submit bids to "win" a good. This is *not* the case here, and we briefly elucidate some of the finer nuances of our application at a higher level.

The VCG mechanism can be applied to any setting with a set of mutually exclusive outcomes experienced by a set of agents. In

Manuscript received 29 June 2021; revised 11 November 2021; accepted 4 January 2022. Date of publication 24 March 2022; date of current version 9 December 2022. This work was supported in part by the National Robotics Initiative 2.0 under Grant 2021-67021-33449 and in part by the Singapore's National Research Foundation as part of the Singapore-Berkeley Building Efficiency and Sustainability in the Tropics (SinBerBEST) program, within the CREATE Program and under the Berkeley Educational Alliance for Research in Singapore (BEARS). *(Corresponding author: Yashaswini Murthy.)*

Ioannis Konstantakopoulos is with Amazon.com, Seattle, WA 98109 USA (e-mail: ioanniskon@berkeley.edu).

Kristy Hamilton is with the Department of Communication, University of California, Santa Barbara, CA 94551 USA (e-mail: kristyhamilton@ucsb.edu).

Yashaswini Murthy, Tanya Veeravalli, and Roy Dong are with the Department of Electrical and Computer Engineering, University of Illinois, Urbana-Champaign, IL 61820 USA (e-mail: ymurthy2@illinois.edu; veerava2@illinois.edu; roydong@illinois.edu).

Costas Spanos is with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94551 USA (e-mail: spanos@berkeley.edu).

This work involved human subjects in its research. Approval of all ethical and experimental procedures and protocols was granted by the UC Berkeley Committee for Protection of Human Subjects under Protocol No. 2013-06-5363.

Digital Object Identifier 10.1109/JSYST.2022.3142443

auctions, the different, mutually exclusive outcomes are which bidder wins the good (there can be a single winner amongst the bidders). In our application, the outcomes are the three possible light settings. In its full generality, the VCG mechanism simply takes into account user preferences and outputs a socially optimal outcome after having reconciled these reported preferences. In smartSDH, users bid for an optimal light setting. The chosen light setting may not be the optimal light setting for every user, yet, due to the shared nature of the office space, only one light setting can be chosen (i.e., the three outcomes are mutually exclusive) and every user experiences the effect of the outcome. Depending on the socially optimal outcome, the mechanism then issues payments to users based on the preferences of all other users, incentivizing truthful reporting of their preferences.

While the VCG mechanism has many desirable properties when all participants act as rational agents, the literature of real-world implementations outside the context of auctions is limited. Therefore, the purpose of this experiment is to 1) assess the feasibility of the modified VCG mechanism in the context of smart building control in group settings and 2) evaluate the effect of smartSDH on light level satisfaction, incentive satisfaction, and energy consumption across time. Furthermore, due to the application domain, minor modifications were required to implement VCG mechanisms in a real-time setting where users cannot walk away and the outside option is more complicated than "nonparticipation."

The rest of this article is organized as follows. In Section II, we discuss the existing literature most relevant to our work and contextualize our contribution. In Section III, we present the mathematical formalism for modeling the decision-making processes of users and the formulation of the mechanism implemented by smartSDH. In Section IV, we outline our experimental setup. In Section V, we present the statistical analysis of our experimental data. In Sections VI and VII, we discuss and interpret the results as well as provide closing remarks summarizing our discoveries and potential avenues for future work.

II. Background and Related Work

In this section, we discuss the existing literature and works most closely related to ours. Broadly speaking, this literature can be broken down into two categories: works on building control and works on mechanism design.

A. Building Control

Nearly 50% of the energy consumed in the US is accounted for by residential and commercial buildings [5], and well-designed algorithms for the control of lighting and heating–ventilation–air-conditioning (HVAC) systems in buildings promise significant benefits for the stability and efficiency of our power grid. Much of the research in building control focuses on finding algorithms that take into account the uncertainty about external factors and operating conditions, while satisfying the constraints introduced by user preferences [6], [7].

Recent work in smart building infrastructures incorporate occupant preferences about thermal comfort [8], satisfaction/well-being [9], lighting comfort [10], acoustical quality [11], indoor air quality [12], indoor environmental monitoring [13], privacy [14], and productivity [15], while simultaneously optimizing for energy efficiency and agile connectivity with the grid. Alongside these, ontological modeling has emerged as one of the latest avenues pursued in the domain of resource allocation through smart technologies [16], [17].

The most relevant work closely related to the experimental study presented in this article concerns the theoretical study of implementing the VCG mechanism for the purpose of moderating the users' thermal comfort across multiple zones within a building [18]. The primary approach of [18] involves modeling of the users' comfort as a function and incentivizing the users to aid in truthful representation of their preferences. Although their simulations and theoretical properties support the utility of such an approach in real life, it lacks strong experimental validation. We aim to bridge this gap by presenting a thorough experimental study of such a mechanism in real world. Consequently, our results indicate the need to develop new mechanisms as the user satisfaction over time in reality does not align with the theoretical predictions of the VCG mechanism.

Prior work has attempted to correct the information imbalance by modeling occupant preferences by the means of a multiagent system [19]. An initial control policy is created from both the building manager's and the occupants' preferences. Then, a rule engine finds a compromise in the system and iteratively performs a compromise. In another related work, [20] demonstrates an auction-based apparatus where each room in an office building makes a bid between the desired temperature and the actual temperature.

With regard to other building control scenarios, [21] carries out an experiment with sensors placed in buildings and uses machine learning and statistical learning techniques with the information from the sensors to diagnose building operation problems related to energy usage and occupant comfort. Along a similar vein, [22] develops a methodology using causal learning statistics to model occupants' thermal discomfort in smart buildings using temperature sensors. A more classical approach to understanding HVAC settings in a building using optimal control is explored in [23]. Here, a variety of optimal control and optimization methods are analyzed in the context of energy efficiency and/or cost-efficient control of HVAC systems, many of which utilize variations of different mechanisms.

A common goal for those implementing building control is energy efficiency. The work [24] proposes a smart software system that balances the requests from different stakeholders (owners, operators, etc.) and also considers the preferences and locations of individual occupants. The framework enables building control while reducing energy consumption and maximizing occupant comfort.

In contrast to the above methods for human–building interaction, we present an experimental study of a different approach. Although some theoretical works consider the issue of incentive compatibility, the experimental studies all primarily assume that participants will truthfully report their preferences *a priori*. Our work assumes that participants will report in their own best interest. As noted in Section I, this may not always coincide with truthful reporting. At the core of our approach is the design of a mechanism that can find the shared conditions (e.g., lighting and HVAC) that fairly account for all occupants' preferences and provide rewards to those who are willing to compromise on the shared conditions for their coworkers. Additionally, our mechanism can pass on some of the building manager's energy-efficiency incentives onto the building occupants so that they can also experience the benefit of energy-efficiency programs. Our framework is centered around a real-time application of a modified VCG mechanism; occupants are asked for their preferences on the shared source, and this information is used to calculate the social optimum and allocate rewards. The rewards are calculated such that occupants can only benefit from the mechanism, and none of the occupants have any incentive to misrepresent their preferences.

B. *Mechanism Design*

Game-theoretic models have been widely used to model selfish agent behavior in a wide variety of applications, ranging from traffic network flow allocation to smart grid optimization. There is one thing common in the core of these problems: each user is trying to maximize a personal utility function (which can be function of various internal and external factors) subject to personal constraints. Mechanism design is an approach to intelligently design incentives for these users toward a common objective. Most game-theoretic analyses rely on the assumption that the utility functions are known *a priori*, which may not be a realistic assumption to make in real-world situations; this is especially true in the case of the energy industry where there are many "noise" variables affecting how well a building can operate. However, there have been some attempts at user preference modeling in intelligent environments [25], some delving into argumentation methods in order to formalize commonsense reasoning [26].

While VCG mechanisms have many nice theoretical properties, it is rare to see their supposed effect in the real-world integrated with decision-making. There are some existing works in the body of literature that contribute to the simulation and/or deployment of VCG mechanisms [18]. For instance, the mechanism is applied to numerical examples of wholesale energy markets to achieve social optimality by incentivizing truth-telling [27]. Similarly, there exists compelling work on simulating the VCG mechanism results in optimal energy load scheduling where social welfare is also measured monetarily [28]. With regard to human computation, crowdsourcing tasks and designing optimal pricing policies also require incentive-compatible mechanism design; strong theoretical guarantees are shown with regard to case studies done using Mechanical Turk [29].

We would like to emphasize that to the best of our knowledge, our work is the first to use VCG mechanisms in a real-world building control application. It is also one of the first papers to implement VCG mechanisms in real-world settings outside of the classical domains of auctions. We hope that this literature develops more in the coming years, as new technologies will have to reconcile the opinions of many users, and an understanding of how humans interact with different mechanisms is crucial for these systems to operate as desired.

III. Model

In this section, we briefly introduce the mathematical formalism for modeling the decision-making processes of users as well as the formulation of our implemented mechanism. To recap, every user submits his/her bid for an optimal light setting, and the mechanism computes the socially best light intensity although it need not be optimal for every individual user. In later sections, we will use data from our experiments to explore the faithfulness of these models.

A. *User Model*

First, we formally outline the assumptions of our human models. We begin by introducing some notation. Let $\mathcal{I}$ denote the finite set of participants, and let $\mathcal{X}$ denote the finite set of outcomes. Recall that our mechanism will choose one of the outcomes in $\mathcal{X}$ to implement. Let $n = |\mathcal{I}|$ denote the number of participants and $m = |\mathcal{X}|$ denote the number of outcomes. Without any loss of generality, we will let the labels of users be $\mathcal{I} = \{1, 2, \ldots, n\}$.

In our smartSDH experiment, the number of outcomes $m = 3$ and corresponds to three different light settings: ("Normal," "Bright," and "Very Bright"), and we have $n = 27$ participants.

In order to impart a mathematical structure to our model, it is necessary to define certain notions.

Definition 3.1 (Type of the user): For each user $i \in \mathcal{I}$ and outcome $x \in \mathcal{X}$, let $\lambda_x^i \in \mathbb{R}$ denote the cost user i incurs when the chosen outcome is x. We refer to the vector $\lambda^i = (\lambda_x^i)_{x\in\mathcal{X}} \in \mathbb{R}^m$ as the *type* of the user.

Our mechanism will choose an outcome x and issue each user $i \in \mathcal{I}$ a payment p_i. We wish to model people's preferences not only across outcomes x but also across outcome-incentive bundles (x, p). This allows us to compensate users for compromising on the outcome; for example, if we choose an outcome x that is particularly distasteful to a user $i \in \mathcal{I}$ (i.e., λ_x^i is very large), then we can compensate by giving them a very large gift card (i.e., choosing p_i to be very large as well). This concept is better illustrated as the utility function of the user.

Definition 3.2 (Utility function): Each user $i \in \mathcal{I}$ makes decisions to maximize their utility function, which is parameterized by their type λ^i and takes as inputs an outcome-incentive bundle (x, p)

$$u(x, p; \lambda^i) = p - \lambda_x^i. \tag{1}$$

The interpretation is that when a user's type is λ^i, the chosen outcome is x, the incentive payment received is p, and their utility is $u(x, p; \lambda^i)$. We note that the utility function is completely specified given the user's type λ^i. Furthermore, since there are only finitely many outcomes, the form of (1) encapsulates any utility function which is quasi-linear in the payments.

We are now ready to introduce our assumptions on our human decision-making model.

Assumption 3.3 (User model): The user's type λ^i captures all the relevant information for their decision-making process and the utility function is given by 1.

Additionally, we will use the notation $\lambda = (\lambda^1, \ldots, \lambda^n)$ to denote the types of *all* users. As per common game-theory convention, we will use the notation $\lambda^{-i} = (\lambda^1, \ldots, \lambda^{i-1}, \lambda^{i+1}, \ldots, \lambda^n)$ to denote the types of all users other than user i. Hence, $f(\lambda) = f(\lambda^i, \lambda^{-i})$.

Assumption 3.4 (Informational structure): We assume that each user $i \in \mathcal{I}$ knows their own type λ^i. Additionally, no one other than i knows λ^i. This includes other users $j \neq i$ and the mechanism designer.

Assumption 3.4 sits at the core of what we wish to model. When asked about their preferences, we assume users will answer based on their desired outcome, not their true preferences. Similar to the motivating example in Section I, where the user suggests a wrong preference in an averaging mechanism, to obtain a final result of his preference, he might report a wrong type with skewed costs as well to bias the VCG output in his favor. Our goal is to design a mechanism that can calculate the outcome that is most socially desirable for all participants *without* access to information about the user types λ. In particular, our mechanism will ask users to report their types; we let their reported values be denoted as $\widehat{\lambda}$. Our mechanism must decide payments and the outcome based only on these reported values $\widehat{\lambda}$.

B. Vickrey–Clarke–Groves Mechanisms

Given the user models, we can analyze whether mechanisms will achieve desired outcomes. In this subsection, we will define VCG mechanisms, outline the desirable properties these mechanisms have in theory, and discuss extensions required for our application.

First, let us define the socially desirable outcome.

Definition III.5 (Social welfare): The social welfare of an outcome $x \in \mathcal{X}$ for a set of users of type $\lambda = (\lambda^1, \ldots, \lambda^n)$ is given by

$$s(x, \lambda) = -\sum_{i \in \mathcal{I}} \lambda_x^i. \tag{2}$$

We define the social optima as the maximizers of the social welfare

$$s^*(\lambda) = \arg\max_{x \in \mathcal{X}} s(x, \lambda). \tag{3}$$

Intuitively, the social welfare function calculates the sum total of everyone's utilities, excluding the incentive payments. In other words, the social welfare function assesses how good each outcome $x \in \mathcal{X}$ is for all participants without payment compensation.

Now, we will provide a definition of a class of mechanisms known as VCG mechanisms and outline some of their theoretically desirable properties as well as nuances in our experiment's particular implementation.

Definition III.6 (VCG mechanisms): Given the reported types $\widehat{\lambda}$, a *VCG mechanism* chooses the outcome that maximizes the social welfare

$$f(\widehat{\lambda}) = \arg\max_{x \in \mathcal{X}} s(x, \widehat{\lambda}). \tag{4}$$

Additionally, for each user $i \in \mathcal{I}$, it pays

$$p_i(\widehat{\lambda}) = h_i(\widehat{\lambda}^{-i}) - \sum_{j \neq i} \widehat{\lambda}^j_{f(\widehat{\lambda})}. \tag{5}$$

Here, h_i is an arbitrary function that does not depend on the reported value $\widehat{\lambda}^i$ of user i. The form of h_i determines the individual optimality of the chosen output with respect to an output which is obtained in the absence of such a mechanism. Depending on the modeling of utilities in various scenarios and constraints over the types of users, h_i can be wisely chosen to ensure the optimality of the VCG mechanism.

VCG mechanisms choose the outcome that is socially optimal for the reported types. Furthermore, it issues payments to every user i based on the utility of all other users from the chosen outcome. Intuitively, this means that once the payments are incorporated, the utility function of every user essentially becomes the social welfare function along with payment compensation. Formally, if we plug in p_i to the utility function of user i, we get

$$u(f(\widehat{\lambda}), p_i(\widehat{\lambda}); \lambda^i) = h_i(\widehat{\lambda}^{-i}) - \sum_{i \in \mathcal{I}} \widehat{\lambda}^i_{f(\widehat{\lambda})} = h_i(\widehat{\lambda}^{-i}) + s(x, \lambda). \tag{6}$$

We next introduce the notion of incentive compatibility in order to incentivize truthful reporting of preferences. This is an important notion in the implementation of the VCG mechanism as it sets the rationale for the users to not lie about their types in order to maximize their utility.

Proposition 3.7 (Incentive compatibility): For every set of user types λ, and for every user i and reported type $\widehat{\lambda}^i$, we have

$$u(f(\lambda), p_i(\lambda); \lambda^i) \geq u(f(\widehat{\lambda}^i, \lambda^{-i}), p_i(\widehat{\lambda}^i, \lambda^{-i}); \lambda^i). \tag{7}$$

The proof for Proposition 3.7 can be found in [4].

Incentive compatibility should be interpreted as follows: λ denotes the true types of all the users, and $\hat{\lambda^i}$ is a possible lie which user i could report. The left-hand side of (7) is the utility user i gains when everyone is truthful. The right-hand side is the utility i derives when everyone except i truthfully responds. Equation (7) states that any such deviations will only lower user i's utility. Thus, truthful reporting by all users forms a Nash equilibrium. Note that this utility function is still parameterized by the true type λ^i as this determines the utility the user actually experiences.

Next, we cover another desirable property for VCG mechanisms: individual rationality. In the typical applications of the VCG mechanism, users can opt-out and choose an outside option. For example, in auctions, a participant can simply choose to not bid at all and walk away. They will receive no utility from receiving a good, and they will lose no utility from paying the auctioneer.

However, we note that in our smartSDH experiment, this is not very well defined. Even if a participant does not report a vote to our mechanism, they must sit in their office space and experience the light chosen by our mechanism. This warrants a minor modification of the typical individual rationality constraint since the outside option becomes type-dependent, which we outline here. To understand this deviation from original VCG mechanism, we introduce the notion of nominal outcome which is analogous to the "outside option."

Assumption III.8 (Nominal outcome): Without any mechanism in place, there is a *nominal outcome* $x_0 \in \mathcal{X}$ that would occur. For each user, their *outside option* is $u(x_0, 0; \lambda^i)$, which is the utility they get from the nominal outcome and no awarded points. It is important to note that $\lambda^i_{x_0}$ need not necessarily be equal to zero for all users. The modified VCG mechanism yields an output which is individually optimal with respect to the nominal outcome, if it satisfies ex-post individual rationality. A VCG mechanism has *ex-post individual rationality* if for all user types λ and any user i

$$u(f(\lambda), p_i(\lambda); \lambda^i) \geq u(x_0, 0; \lambda^i). \tag{8}$$

The interpretation is that the nominal outcome (outside option) would be the chosen outcome in the absence of our mechanism. When we enforce ex-post individual rationality, this means that every user is better off when the mechanism is implemented, as opposed to when the mechanism does not exist at all.

It is important to note that this is a deviation from typical applications of the VCG mechanism. In smartSDH, a user cannot "walk away" from their shared office space; rather, we have to ensure that their utility increases as a result of our mechanism's deployment. From a technical perspective, the important distinction is that the utility of the outside option depends on the user's type; this is what necessitates modification at a theoretical level.

Assumption III.9 (Bounded user utilities): The types of all users are bounded between 0 and $\lambda_{\max}$, i.e., $\lambda^i_x \in [0, \lambda_{\max}]$ for all $i \in$ and $x \in \mathcal{X}$.

Proposition III.10 (Existence of an ex-post individually rational mechanism): Using $h_i(\widehat{\lambda}^{-i}) = n\lambda_{\max}$, the VCG mechanism is ex-post individually rational.

Proof: First, note that $n\lambda_{\max} - \sum_{i\in\mathcal{I}} \lambda^i_x \geq 0$ for any x and any λ. Thus, $u(f(\lambda), p_i(\lambda); \lambda^i) \geq 0$. Next, note that since $\lambda^i_{x_0} \geq 0$ for all i by assumption, we have $u(x_0, 0; \lambda^i) \leq 0$. The inequality in (8) follows.

In the smartSDH experiment, $h_i(\widehat{\lambda}^i) = n\lambda_{\max}$ happens to be independent of the types of the users. It can be thought of as the basic wage that every user is given for participating in the study.

Lastly, we outline the last modification needed to implement a VCG mechanism in our office space setting. Users can come and go at any point in time, and they can modify their reported preferences at any point in time as well. The modification needed from the classical VCG mechanism applications is the temporal aspect of this.

Essentially, this modification is quite simple: transform all the quantities discussed above into rates. Rather than interpreting λ^i_x as the cost user i incurs when the outcome is x, we interpret λ^i_x as the cost per hour user i incurs. Similarly, if the VCG mechanism decides to pay user i a payment of p_i, they are paid p_i points per hour. Then, any time a user enters, leaves, or modifies their vote, we treat that segment as one round of the mechanism, weighted by its duration.

More formally, suppose the set of users logged on and their reported preferences $\widehat{\lambda}$ are constant on the time interval $[t_0, t_1]$. Then, in that time interval, our mechanism chooses the outcome $f(\widehat{\lambda})$. Each user i who is logged on in that time interval receives a reward $(t_1 - t_0)p_i(\widehat{\lambda})$ for their participation during the time interval $[t_0, t_1]$.

C. Contributions and Goal of the Investigation

The goal of the present investigation was to test the application of the modified VCG mechanism to a smart building control application, dubbed "smartSDH" and to experimentally verify its utility. As mentioned previously, a VCG mechanism selects the socially optimal outcome among a set of possible outcomes—in this case, the preferred brightness of overhead office lights—and then issues payments to users based on the decisions of all other users. In this sense, it is a mechanism where being truthful is the best strategy for the individual and the group. Although VCG mechanisms have many desirable theoretical properties when all participants act as rational agents, real-world applications of VCG mechanisms are limited.

This study analyzes perceptions and behaviors of users in a shared office space who interacted with a VCG-operated smartSDH over a five-month period. Participants interacted with smartSDH via a web portal, where they reported their preference for the brightness of the overhead office lights in real time so that the VCG mechanism could determine the socially optimal light setting for the group. Users were allotted points according to the modified VCG mechanism. Whenever the total points earned crossed a threshold, a lottery was held for gift cards; over the five-month period, we rewarded a total of $2900 worth in gift cards. Additionally, to emphasize the noncompetitive aspect of sharing an office space, whenever the total points earned crossed another threshold, we hosted a catered lunch for all the participants of smartSDH. Because adaptation to technology often evolves over time [30], sometimes gradually and sometimes sporadically, we chose to evaluate the behavioral outcomes over the course of three distinct time periods (T1 = Wk1–Wk7; T2 = Wk8–Wk15; T3 = Wk16–Wk 22). Upon observation of data, the temporal changes were significant across periods of eight weeks; hence, the analysis breaks down the time period of data analysis to eight weeks with no gaps in between. Whether or not this duration is a a generalizable phenomena (e.g., if people tend to always change opinions after roughly eight weeks with a new technology) or whether the demographics of our participants influenced this duration are interesting questions for future study. In such a setting, three research questions become salient.

- RQ1: What is the influence of using the VCG-operated smartSDH on light and incentive satisfaction across time? How does an average user's lighting preference vary across time and in response to the incentives?
- RQ2: What is the influence of using the VCG-operated smartSDH on energy saving?
- RQ3: Is it possible to determine user's preferences from their immediate environment to facilitate smart lighting without user intervention? What is the relationship between light level preferences and atmospheric conditions including humidity, temperature, pressure, and solar radiation?

IV. Method

A. Participants

Twenty-seven graduate students from a public research university in California, USA participated in the longitudinal study in exchange for prize items, which were allotted based on their performance in the point-based system. The preliminary survey was optional, and out of the 27 participants, 18 responded. The survey asked the students about what they considered to be their ideal incentives (gift cards vs. fitbits, etc.), the duration of their commute, modes of transportation they resorted to, time of arrival to the office, etc. The majority of users were men (gender: 73% men; 22% women; 5% prefer not to respond) with 50% of users between the ages of 22 and 25, while the rest were above 26 years old. The majority of users had incomes between \$35 000 and \$40 000, and the rest had lower incomes. Before commencing the experiment, users were somewhat satisfied with the light conditions in the office ($M = 4.39$, SD $= 1.14$) on a 5-point scale. As for ideal incentives, nearly everyone preferred gift cards from Amazon, iTunes, and Google Play (94%) over the other options, which were complimentary vouchers for drinks at a campus coffee shop or lotteries for big-ticket items such as Apple Watches, Fitbits, and EarPods.

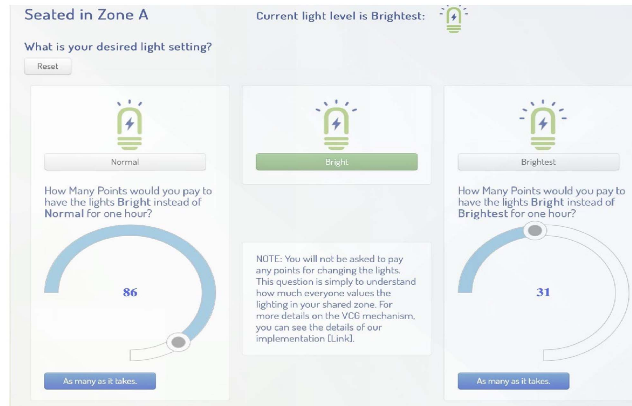

Fig. 1. Voting portal for smartSDH, where users log in and report their preferences on how many points they would pay to change a light setting from the currently implemented setting.

B. Procedure

Participants in the longitudinal study interacted with smartSDH in an open office space with cubicles on a university campus over a period of five months. The desks in the office space were divided into different lighting zones with a set of overhead lights serving as the primary source of illumination for each zone. The smartSDH operated independently for each zone. From the work hours of 9 A.M. to 5 P.M. local time, participants were only able to control their office lights through the smartSDH web portal. After work hours, the light switches returned to normal operation.

Each participant had access to the password-protected smartSDH web portal, visualized in Fig. 1. In addition to allowing users to vote for their light settings in real time, the web portal also provided participants with visualizations of the state of the office space. Users could view their personal point totals, the currently implemented light setting, and the global progress to the individual and communal incentive thresholds. They could also see which occupants were present in their zone. Users were able to monitor their floor's lighting lumen level and temperature in real time, with a refresh interval of 1 s. A view of the designed portal can be seen in Section IV-B. To adjust the light conditions, participants first selected their preferred light setting from three available options, "Normal," "Bright," and "Very Bright," which corresponds to 33%, 67%, and 100% of the maximum possible illumination. We measured the illumination at a typical desk under these settings using the Konica Minolta T-10 A illumination meter at night. The "Normal (33%)" setting corresponds to 122 lx, "Bright (67%)" setting corresponds to 215 lx, and "Very Bright (100%)" to 431 lx. These readings were captured without any other light sources at night and represents the contribution to illumination from the overhead lights.

Once participants selected their preferred light setting (e.g., "Bright"), the web portal requested follow-up information about the two settings not chosen (e.g., "Normal" and "Very Bright"). Specifically, the web portal asked participants to indicate the extent to which they preferred their chosen light setting over each of the alternative settings (e.g., "How many points are you willing to pay to have the lights set to BRIGHT instead of NORMAL for one hour?") from 0 to 100 points (see Section IV-B for illustrations of the voting procedure). These two measures—preferred light setting and relative preference—contributed to the VCG mechanism's chosen outcome. Every time a user logged on, logged off, or changed their reported preference, smartSDH calculated $f(\hat{\lambda})$ based on the reported values of all users who were currently logged on and rewards each user i's account with points $p_i(\hat{\lambda})$, according to the calculations in Section III. Users were able to modify their reported preferences at any time; however, they were required to be present in the office to vote. Office presence was enforced via the browser's geolocation data.

In addition to the points earned from the VCG mechanism, we also rewarded users for completing a repeatable survey about their experiences with smartSDH. This repeatable survey included Likert-scale questions about the incentives provided, the design of the web portal, the level of comfort participants experienced, their satisfaction with the current light setting, the level of awareness about energy-saving actions they could take, and their productivity in the office.

C. Rewards

The points were converted to values by the users in two ways throughout the study. First, whenever the total points earned by all participants exceeded a threshold, we held a lottery for gift cards. The probability of one winning the lottery was proportional to the number of points in one's account, and multiple gift cards were given each time the lottery threshold was reached. This was the individual incentive. Second, their points built toward a communal incentive. Whenever the total points earned by all participants exceeded the communal incentive threshold, we provided a catered lunch to participants.

D. Apparatus

The users interacted with a web portal where they submitted their preferences for the overhead lighting in the office. The

Building Automation and Control (BACnet) protocol was used for communications between the web portal interface and the light actuators in the space. We used BACNet because it is the prominent protocol for HVAC applications.

In practical scenarios, the VCG mechanism described here would easily translate to real-world deployments, so long as the building was equipped with some protocol by which to control the lights and interface with participants. The computations of the mechanism are very simple to implement once the reported preferences are collected.

E. Measures

Light setting preference: Participants selected one of three available settings: ("Normal," "Bright," and "Very Bright"), which corresponds to 33%, 67%, and 100% of the maximum possible illumination.

Relative preference: For the two settings not chosen, the web portal asked "How many points are you willing to pay to have the lights set to PREFERRED instead of ALTERNATIVE for one hour?" For each of the two alternatives, users provided an integer between 0 and 100. We assumed $\lambda_{\max} = 100$, where $\lambda_{\max}$ is as defined in 3.9. As a quality-of-life feature, we also included a button that allowed users to set their vote to $\lambda_{\max}$ with one click.

Light level satisfaction: Two items measured participants' satisfaction with their lighting conditions on a given day, "I am satisfied with today's lighting conditions" and "Today's lighting conditions were uncomfortable." Responses were recorded on a 5-point Likert-type scale from 1 (strongly disagree) to 5 (strongly agree) ($\alpha = 0.91$).

Incentive satisfaction: Two items measured participants' satisfaction with the incentives provided on a given day, "I am happy with the current incentives provided" and "The current incentives are not satisfactory." Responses were recorded on a 5-point Likert-type scale from 1 (strongly disagree) to 5 (strongly agree) ($\alpha = 0.86$).

User interface satisfaction: Two items measured participants' satisfaction with the web portal used to manipulate the smart lights on a given day, "I am satisfied with the current web interface" and "The web portal leaves much to be desired." Responses were recorded on a 5-point Likert-type scale from 1 (strongly disagree) to 5 (strongly agree) ($\alpha = 0.78$).

Energy consumption (% savings/time): We measured energy consumption in terms of percentage savings. That is, we calculated the implemented lighting over a baseline of 100% lighting. This is discussed in Section V-B.

Humidity (%): Atmospheric sensors in the office space measured relative humidity in percentage water vapor—where 100% corresponds to fully saturated air at dewpoint. Atmospheric sensors placed in the room.

Temperature (°F): The above-mentioned sensors measured the temperature. Another atmospheric condition captured by the aforementioned sensors.

Pressure (Hg): The sensors measured barometric pressure in terms of units of mercury (Hg). More precisely, a unit of Hg denotes the pressure exerted by a column of mercury 1 in (25.4 mm) in height at the standard acceleration of gravity. Another atmospheric condition captured.

Solar radiation (W/m^2): Also known as solar irradiance, and it is measured in terms of power per unit area (Watts per square meter, in this case).

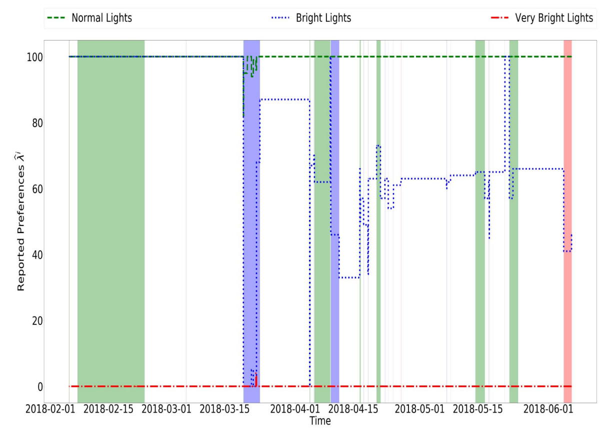

Fig. 2. Voting history for a player across all lighting options for the entire gaming period. The unshaded rectangular regions in the background represent the player's absence. The green region indicates normal setting, purple indicates bright, and red indicates very bright setting.

V. Results

Fig. 2 shows the voting behavior of a user throughout the five-month period of the game. The *y*-axis, Reported Preferences, is the number of points the user is willing to pay to change the lights when it is at that setting. For instance, in this case, the user is willing to pay nothing to change the "Very Bright Lights" setting because it is their most preferred setting and would prefer that it remains at that setting. The shaded regions in the background show the implemented light setting during that time; blank regions indicate that the user was not logged in at the time or actively voting. In this instance, the user initially signals that they do not favor the dimmest light settings and medium light settings and the VCG mechanism implements the dimmest light. As time evolves, the user realizes they are more willing to accept the medium light setting since compromise has proven to output a brighter light overall, as seen in the short instances of blue and, eventually, red shading in the background.

A. Effect of smartSDH on Light Level and Incentive Satisfaction Across Time (RQ1)

Satisfaction with the light levels across time while controlling for user interface satisfaction was analyzed using a three (T1 vs. T2. vs. T3) within-subjects analysis of variance (ANOVA). We observed a significant effect of time on light level satisfaction after controlling for user interface satisfaction, $F(2,151) = 4.21$, $p = 0.017$, $d = 0.217$. *Post hoc* information reveals that participants reported greater satisfaction with light levels at T1 $(M = 3.44$, SD $= 1.03)$ than T2 $(M = 2.94$, SD $= 1.11$, $p = 0.039)$ and T3 $(M = 2.96$, SD $= 1.31$, $p = 0.018)$. These results suggest that after an initial period of satisfaction with

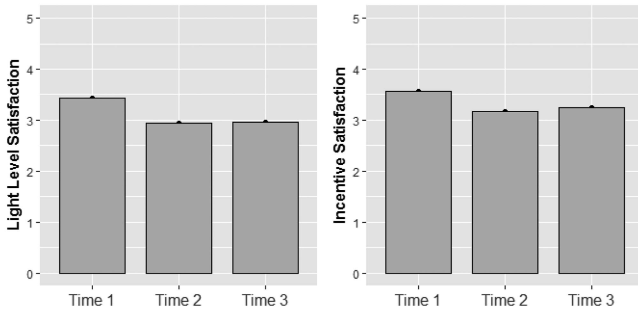

Fig. 3. Effect of smartSDH on light level and incentive satisfaction across time ($N = 157$).

the light levels, participants became less satisfied with light levels (Fig. 3).

Satisfaction with incentives across time while controlling for user interface satisfaction also was analyzed using a three (T1 vs. T2. vs. T3) within-subjects ANOVA. We observed a significant effect of time on incentive satisfaction after controlling for user interface satisfaction, $F(2, 151) = 3.70, p = 0.027, d = 0.171$. *Post hoc* information reveals that participants reported greater satisfaction with incentives at T1 $(M = 3.57$, SD $= 0.82)$ than T2 $(M = 3.16,$ SD $= 1.11,$ $p = 0.001)$ and T3 $(M = 3.24,$ SD $= 1.37, p = 0.001)$. Levene's test indicated that the assumptions of homogeneity in variance were met for the effect on light level satisfaction ($p = 0.509$) but not for incentive satisfaction ($p = 0.028$). We cannot necessarily rely on the test of difference regarding incentive satisfaction because the basic model violates the assumption of equal population variance. Still, it seems reasonable to conclude that despite the theoretical properties of VCG mechanisms, the experimental data suggest the VCG mechanism does not have an appreciable effect on light level and incentive satisfaction of users, as shown in Fig. 3.

B. Energy Saving During smartSDH Use Over five-Month Period (RQ2)

To evaluate the effect of the VCG-operated smartSDH on energy consumption, we compared the intensity of the overhead lights in the office space over the five-month period of smartSDH use to the intensity of the overhead light during normal operation (i.e., 100% of the maximum possible illumination). Our results indicate that employing smartSDH in the office from 9 A.M. to 5 P.M. reduced energy consumption by 35.22% over the five-month study period.

It is of interest to note that these are energy savings which are achieved *without* any incentives promoting energy efficiency. Prior to the deployment of our experiment, this office space exclusively used the 100% light setting. The energy savings here are achieved from participants stating they *preferred* the overhead lights be dimmer, independent of any incentives. In other words, we can view this as energy savings resulting from conserving energy that participants did not want to use. Directly incentivizing energy efficiency is a direction for future research, which we discuss in Section VI.

TABLE I
DESCRIPTIVE STATISTICS AND CORRELATIONS FOR ATMOSPHERIC VARIABLES ON REPORTED LIGHT LEVEL PREFERENCES

	Mean	*SD*	*r*	*p*
Light setting preference	61.81	27.84	1.00	
Temperature (deg F)	55.00	5.17	.017	.005
Pressure (Hg)	29.62	0.11	.06	.292
Solar radiation (W/m^2)	249.21	370.72	.27	.001
Humidity (%)	79.93	11.05	-.18	.003

Note: r represents the zero-order correlation with light level preference. p represents the corresponding p-value for each correlation. $N = 276$.

C. Relation Between Reported Light Setting Preference and Atmospheric Conditions (RQ3)

Table I presents means, standard deviations, and bivariate correlations for atmospheric conditions and reported light setting preference. Light setting preference was significantly and positively correlated with temperature and solar radiation such that participants prefer brighter lights as temperature and solar radiation increases ($p < 0.05$). Light setting preference was significantly and negatively correlated with humidity such that participants prefer brighter lights as humidity decreases ($p < 0.05$). Pressure did not significantly correlate with reported light setting preference. We suspect that this is due, in part, to the fact that temperature, solar radiation, and humidity are correlated with the natural illumination provided, whereas the atmosphere pressure itself directly does not. We take these results to suggest that externally observable factors contribute to light setting preference. Future mechanisms for building control may wish to consider the contribution of real-world observable factors to individuals' light setting preference and reactance to VCG-determined lighting decisions.

VI. DISCUSSION

In this study, we implemented a modified VCG mechanism, smartSDH, to determine the brightness of overhead lights in a shared office space. The goal of smartSDH was to determine the light setting that maximized the sum of everyone's utility fairly. To do so, smartSDH would issue incentives as needed to promote truthful reporting of preferences as well as ensure that all users were better off with smartSDH than a nominal light setting along with no rewards. To this end, VCG mechanisms were an appropriate choice for satisfying many of our desiderata.

Our first research question dealt with the influence of smartSDH on light and incentive satisfaction. In theory, smartSDH should improve the satisfaction of participants with the lighting. Regarding the first hypothesis on measuring the influence of time on light and incentive satisfaction, quantitative results shown in Section V suggest that there is no appreciable increase in satisfaction. We see that in T1, there was greater satisfaction, on average, across all participants than compared to time periods T2 and T3. Additionally, participants were significantly more satisfied with incentives during T1; however, this satisfaction deteriorated in periods T2 and T3. This could be due to a variety of reasons. First, a few survey respondents stated that they were "too busy" to regularly report their preferences, which implies that the user interface for the portal was not as convenient as desired. So as time went on, people felt the task more burdensome resulting in a more sporadic voting pattern as time went on.

Second, there is evidence in the literature about "technology burnouts" (see [31] and the references therein). These references discuss the issues with technology adoption, highlighting the fact that with the right system in place, it is possible for behavioral approaches such as the one in smartSDH to be much more effective in reducing energy consumption, which can improve the result for our second hypothesis on energy saving. This suggests that there was initially a novelty effect for our users, who were more active in the first time period. In the second and third time periods, users were generally less engaged and motivated to continue voting or, simply put, the users experienced a technology burnout. In general, if the user population was more aware of appliance information (in our case, about the lights' energy consumption), it would facilitate greater energy savings. For future work, this suggests more testing on how the user interface is perceived to obtain better insights for design decisions as well as increasing users' awareness on the mechanics of the sensors. If the technology burnout can be contained or even delayed, this methodology can be adopted for a longer period of time with adequate user satisfaction and savings on energy consumption. In particular, this issue of burnout is not purely technological and has social dimensions as well. A very promising direction for future research is understanding the social and the technological approaches to achieve smarter societies and the interrelations between them [32].

Our second research question dealt with the influence of smartSDH on energy consumption. We found that there is a significant reduction in energy usage. This is noteworthy because though the participants generally felt inconvenienced and dissatisfied with the implemented mechanism, their behaviors ended up effectively reducing overall energy usage. For future directions of the work, it would be of interest to relate user satisfaction with the energy consumption. However, our experimental data cannot provide such analysis, as the surveys were measures of overall satisfaction rather than instantaneous satisfaction (i.e., "how satisfied are you right now?"). Additionally, it might be useful to note that this framework readily allows one to incorporate incentives for energy-efficient behavior. For example, if a building costs $c_i \in \mathbb{R}$ dollars to operate at setting $x_i \in \mathcal{X}$ for 1 h, we can introduce another participant of type $\lambda^0 = (c_1, c_2, \ldots, c_m)$. This term now shows up in each participant's utility function through the VCG mechanism, and if everyone gains more than $c_i - c_j$ in utility from choosing i over j, then the mechanism will choose i over j. Each participant can then internalize the incentives, directly experience its benefits, and, as a group, decide whether or not the difference in cost warrants a change in settings. This would allow users to internalize the externality—in other words, to shift the external cost to their own internally accrued cost. However, for our experiment itself, we did *not* incorporate any energy-efficiency incentives; our primary goal was to find the light settings that were the most preferred by all the occupants in the office space. Moreover, technology adoption rates are generally higher for software solutions such as smart meters (analogous to the sensors used in smartSDH) since the hardware is generally installed by an external party such as the utility company. Since this has virtually no cost to the user, along with zero personal installation effort, there is much to be said about integrating smart software solutions to reduce energy consumption in an aggregate energy source.

Our last research question dealt with the relationship between the preferences of users and atmospheric conditions. Our results show that there is a significant relationship between participants' votes and factors such as solar radiation and humidity. However, there is not much qualitative information that we could obtain about why participants felt that these factors influenced them to vote in a particular manner. These qualitative aspects of the user's voting can be asked in a survey at the end of the experiment, which is future work when conducting such experiments. The data from this experiment suggests that estimating these preferences from environmental factors is not a viable avenue because it indicates that mechanisms for smart building control may require occupants to consistently report their preferences. This is because these preferences cannot be accurately predicted from externally observable factors. In addition, the free-form qualitative survey also suggests that people experienced no appreciable increase in satisfaction as a result of the mechanism. We explore a few potential reasons for this, all of which are interesting avenues for future study.

We note that, due to limitations of this study, our geolocation methods required users to regularly sign in from their browser; otherwise, users could vote remotely or leave their computer on to earn points illicitly. We hope that future studies can explore methods to reduce the intrusiveness of these mechanisms, as real-world products would likely have significant investment in user-experience design. This is outside the scope of our study, however.

Another possible reason for satisfaction levels not being significant might have been because $\lambda_{\max}$, as defined in Assumption 3.9, was set too low. This again was due to the limitations of our study. Our value for $\lambda_{\max}$ was determined by our experimental budget. That is, we set $h_i = n\lambda_{\max}$ as a base rate for participation and needed to bound how much money we would give out in rewards over the duration of the experiment. For smartSDH, our budget was over \$100 per participant, which outweighs the monetary savings from reducing energy consumption. In practical applications, it may be beneficial to relax the requirement that all users are better off due to the mechanism. Under this relaxation, some users may wind up paying money into the mechanism, but it would still implement the socially optimal outcome. This ties back in with the idea of "internalizing the externality" that may be caused due to differences in preferences of lighting. To implement this, there must be a way to force participation against a user's will. For example, in real-world deployments of mechanisms such as ours, occupants may be asked to pay for the energy consumption at their desk. We note that this is not likely to be viable in an academic research experiment. Alongside the above considerations, it is important to note that since VCG is generally employed in auctions, the users involved are assumed to be competitive. According to Assumption 3.4, it is assumed that a user knows their type and others are unaware of it. However, due to the nature of the experiment or the fact that the web interface shows the users in the office environment at any instant, the users might cooperate with one another in determining their types, which might once again have an impact on the optimality of the VCG mechanism-derived outcome.

Finally, we note that many qualitative aspects of our study suggest a status quo effect. Some participants told us during the catered events that they were happy with smartSDH because the lights were typically too bright. In contrast, some other participants insisted that no one would want lights so dim. Anecdotally, it seemed some participants were dissatisfied because they did not believe that the mechanism was implementing the social optimum. It is interesting to note how it is more socially acceptable to walk into a full office and turn up the lights than it is to walk into a full office and turn down the lights; these social contexts likely had a factor in the experience of participants of this study. We think examining the effect of the status quo when new Internet of Things (IoT) technologies are deployed in these settings is a very interesting direction for future research.

VII. Conclusion

In summary, the contributions of this article are as follows.

1) Implementing a modified VCG mechanism to determine the brightness of the overhead lights for 27 participants over a five-month period in their actual office space, giving out $2900 in rewards. We emphasize that it is challenging to implement the VCG mechanism outside of typical auction settings (see Section II-B) and most of the work in the realm of VCG based mechanisms have been strictly theoretical and simulation based.
2) Studying user satisfaction with incentives and light level and the impact of the mechanism on energy savings and consumption.
3) Determining correlations between light setting preferences and atmospheric conditions. Although some factors had significant correlations with the reported preferences, the predictive value was generally very poor. This implies that mechanisms for building control may require users to constantly report their preferences, as it would be difficult to build estimators of these preferences from externally observable factors.
4) Outlining some barriers to the implementation of VCG mechanisms in IoT settings and potential reasons our study did not achieve the expected gains in satisfaction.

Despite the theoretical properties of VCG mechanisms, the experimental data suggests that the VCG mechanism did not have an appreciable effect on the satisfaction or awareness of users in any of these user-perception categories. Our data suggests that rational agent models may require some modifications to capture how humans typically will interact with an IoT technology in the background of their work life. In particular, the preferences reported show some temporal variation that may be due to users learning their own preferences, which is in contrast with the typical rational agent model. Furthermore, there is a need to design the user experience to be as minimally intrusive as possible and potentially relax the requirement that all users are better off with the existence of the mechanism.

VII. Acknowledgments

Ioannis C. Konstantakopoulos was a Scholar of the Alexander S. Onassis Public Benefit Foundation. The authors would like to thank Chris Hsu, the applications programmer at CREST Laboratory, UC Berkeley. The authors would also like to thank Claire Tomlin, without whom this experiment may never have been deployed. Additionally, the would like to thank UC Berkeley Student Technology Fund for the Award of Student Technology Fund Initiative – Large-Scale Project (2018). Finally, the authors would like to thank the participants of smartSDH.

References

[1] S. H. Shah and I. Yaqoob, "A survey: Internet of things (IOT) technologies, applications and challenges," in *Proc. IEEE Smart Energy Grid Eng.*, 2016, pp. 381–385.
[2] W. Vickrey, "Counterspeculation, auctions, competitive sealed tenders," *J. Finance*, vol. 16, no. 1, pp. 8–37, 1961.
[3] T. Groves, "Incentives in teams," *Econometrica*, vol. 41, no. 4, pp. 617–631, 1973.
[4] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2007.
[5] EIA, "Monthly energy review," US Energy Inf. Admin., Washington, DC, USA, 2021. [Online]. Available: https://www.eia.gov/tools/faqs/faq.php?id=86&t=1
[6] Y. Ma, G. Anderson, and F. Borrelli, "A distributed predictive control approach to building temperature regulation," in *Proc. Amer. Control Conf.*, 2011, pp. 2089–2094.
[7] A. Aswani *et al.*, "Identifying models of HVAC systems using semiparametric regression," in *Proc. Amer. Control Conf.*, 2012, pp. 3675–3680.
[8] C. Karmann, S. Schiavon, and E. Arens, "Percentage of commercial buildings showing at least 80% occupant satisfied with their thermal comfort," *UC Berkeley: Center Built Environ.*, 2018. [Online.] Available: https://escholarship.org/uc/item/89m0z34x
[9] M. Frontczak, S. Schiavon, J. Goins, E. Arens, H. Zhang, and P. Wargocki, "Quantitative relationships between occupant satisfaction and satisfaction aspects of indoor environmental quality and building design," *Indoor Air*, vol. 22, pp. 119–131, Apr. 2012.
[10] R. A. Baron, M. S. Rea, and S. G. Daniels, "Effects of indoor lighting (illuminance and spectral distribution) on the performance of cognitive tasks and interpersonal behaviors: The potential mediating role of positive affect," *Motivation Emotion*, vol. 16, no. 1, pp. 1–33, 1992.
[11] E. E. Ryherd and L. M. Wang, "Implications of human performance and perception under tonal noise conditions on indoor noise criteria," *J. Acoustical Soc. Amer.*, vol. 124, no. 1, pp. 218–226, 2008.
[12] J. Sundell *et al.*, "Ventilation rates and health: Multidisciplinary review of the scientific literature," *Indoor Air*, vol. 21, pp. 191–204, Jun. 2011.
[13] M. Jin, S. Liu, S. Schiavon, and C. Spanos, "Automated mobile sensing: Towards high-granularity agile indoor environmental quality monitoring," *Building Environ.*, vol. 127, pp. 268–276, 2018.
[14] R. Jia, R. Dong, S. S. Sastry, and C. J. Spanos, "Privacy-enhanced architecture for occupancy-based HVAC control," in *Proc. 8th Int. Conf. Cyber-Physical Syst.*, New York, NY, USA, 2017, pp. 177–186.
[15] P. Wargocki *et al.*, *REHVA Guidebook No 6 - Indoor Climate and Productivity in Offices - How to Integrate Productivity in Life-Cycle Cost Analysis of Building Services*. Brussels, Belgium: REHVA, 2006.
[16] Z. Nezami, K. Zamanifar, D. Arena, and D. Kiritsis, *Ontology-Based Resource Allocation for Internet of Things*. Berlin, Germany: Springer, 2019, pp. 323–330.
[17] N. Seydoux, K. Drira, N. Hernandez, and T. Monteil, "IoT-O, a core-domain IoT ontology to represent connected devices networks," in *Proc. Knowl. Eng. Knowl. Manage.*, E. Blomqvist, P. Ciancarini, F. Poggi, and F. Vitali, Eds., Cham, Switzerland: Springer International Publishing, 2016, pp. 561–576.
[18] S. K. Gupta, K. Kar, S. Mishra, and J. T. Wen, "Incentive-based mechanism for truthful occupant comfort feedback in human-in-the-loop building thermal management," *IEEE Syst. J.*, vol. 12, no. 4, pp. 3725–3736, Dec. 2018.

[19] M. Boman, P. Davidsson, N. Skarmeas, K. Clark, and R. Gustavsson, "Energy saving and added customer value in intelligent buildings," Inf./Soc./Energy/Syst., Rep., 1970. [Online.] Available: http://www.enersearch.com/company/knowledgebase/publications/by_project/ISES/ises9/paam98/paam98.pdf

[20] S. H. Clearwater, "Auction-based control system for energy resource management in a building," US Patent 5 394 324A, Feb. 1995.

[21] J. Ploennigs and A. Schumann, "From semantic models to cognitive buildings," in *Proc. 31st AAAI Conf. Artif. Intell.*, 2017, pp. 5105–5106.

[22] S. Egedorf and H. R. Shaker, "Adverse condition and critical event prediction in cranfield multiphase flow facility," in *Proc. Int. Conf. Sens. Diagnostics Prognostics Control*, 2017, pp. 557–564.

[23] S. Wang and Z. Ma, "Supervisory and optimal control of building hvac systems: A review," *Heating Ventilation Air Conditioning Res.*, vol. 14, no. 1, pp. 3–32, 2008.

[24] H. Chen, P. Chou, S. Duri, H. Lei, and J. Reason, "The design and implementation of a smart building control system," in *Proc. IEEE Int. Conf. E- Bus. Eng.*, 2009, pp. 255–262.

[25] J. C. Augusto and A. Muãoz, "User preferences in intelligent environments," *Appl. Artif. Intell.*, vol. 33, no. 12, pp. 1069–1091, 2019.

[26] C. Teze, S. Gottifredi, A. Garcia, and G. Simari, "An approach to generalizing the handling of preferences in argumentation-based decision-making systems," *Knowl.-Based Syst.*, vol. 189, 2019, Art. no. 105112.

[27] Y. Xu and S. H. Low, "An efficient and incentive compatible mechanism for wholesale electricity markets," *IEEE Trans. Smart Grid*, vol. 8, no. 1, pp. 128–138, Jan. 2017.

[28] P. Samadi, R. Schober, and V. W. Wong, "Optimal energy consumption scheduling using mechanism design for the future smart grid," in *Proc. IEEE Int. Conf. Smart Grid Commun.*, 2011, pp. 369–374.

[29] G. Goel, A. Nikzad, and A. Singla, "Mechanism design for crowdsourcing markets with heterogeneous tasks," in *Proc. 2nd AAAI Conf. Hum. Comput. Crowdsourcing*, 2014, pp. 77–86.

[30] G. DeSanctis and M. S. Poole, "Capturing the complexity in advanced technology use: Adaptive structuration theory," *Org. Sci.*, vol. 5, pp. 121–147, May 1994.

[31] K. Carrie Armel, A. Gupta, G. Shrimali, and A. Albert, "Is disaggregation the holy grail of energy efficiency? the case of electricity," *Energy Policy*, vol. 52, pp. 213–234, 2013.

[32] C. Vlachokostas, "Smart buildings need smart consumers: The meet-in-the middle approach towards sustainable management of energy sources," *Int. J. Sustain. Energy*, vol. 39, no. 7, pp. 648–658, 2020.

Ioannis C. Konstantakopoulos received the Diploma (Hons.) degree in electrical and computer engineering from the University of Patras, Patras, Greece, in 2012, and the M.S. and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, Berkeley, CA, USA, in 2014 and 2018, respectively.

He was a Scholar of the Alexander S. Onassis Public Benefit Foundation that provides financial support for outstanding Greek Doctoral students studying at U.S. institutions. He is currently an Applied Scientist with Alexa Artificial Intelligence Team, Amazon.com, Inc., Seattle, WA, USA. His research interests include deep learning, statistical learning, machine learning, algorithmic game theory, and optimization.

Kristy A. Hamilton received the Ph.D. degree in communications and media from the Institute of Communication Research, University of Illinois at Urbana-Champaign, Champaign, IL, USA, in 2020.

She is currently an Assistant Professor of Digital Media with the Department of Communication, University of California, Santa Barbara, Santa Barbara, CA, USA. She uses experimental methods from cognitive and social psychology to understand the inherent qualities or liabilities of human memory and cognition in a digital environment. She has authored or coauthored journals such as *Computers in Human Behavior*, *Journal of Applied Research in Memory and Cognition*, *Applied Cognitive Psychology*, and *New Media & Society*.

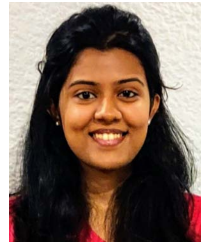

Yashaswini Murthy received the B.Tech. and M.Tech. degrees in mechanical engineering with specialization in computer aided design and automation and a minor in systems and control engineering from the Indian Institute of Technology, Bombay (IITB), Bombay, India, in 2019. She is currently working toward the Ph.D. degree in electrical engineering with the University of Illinois at Urbana-Champaign, Champaign, IL, USA.

Her research interests include domain of control theory.

Tanya Veeravalli received the B.A. degree with majors in computer science and economics from the University of California, Berkeley, Berkeley, CA, USA, in 2019. She is currently working toward the Ph.D. degree in electrical and computer engineering with the University of Illinois at Urbana-Champaign, Champaign, IL, USA.

Her research interests include areas of learning in dynamical systems and optimization.

Costas Spanos (Fellow, IEEE) received the Diploma degree in electrical engineering from the National Technical University of Athens, Athens, Greece, in 1980, and the M.S. and Ph.D. degrees in electrical communication engineering from Carnegie Mellon University, Pittsburgh, PA, USA, in 1981 and 1985, respectively.

In 1988, he joined the Department of Electrical Engineering and Computer Sciences (EECS), University of California, Berkeley (UC Berkeley), Berkeley, CA, USA, where he is currently the Andrew S. Grove Distinguished Professor and the Director with the Center for Information Technology Research in the Interest of Society and the Banatao Institute (CITRIS). He is also the Founding Director and CEO with the Berkeley Education Alliance for Research in Singapore (BEARS), Singapore, and the Lead Investigator of a large research program on smart buildings based in California and Singapore. Prior to that, he has been the Chair of EECS, UC Berkeley, the Associate Dean for Research with the College of Engineering, UC Berkeley, and the Director of the UC Berkeley Microfabrication Laboratory. His research interests include sensing, data analytics, modeling and machine learning, with broad applications in semiconductor technologies, and cyber–physical systems.

Roy Dong received the B.S. (Hons.) degree in computer engineering and the second B.S. (Hons.) degree in economics from Michigan State University, East Lansing, MI, USA, in 2010 and the Ph.D. degree in electrical engineering and computer sciences from the University of California, Berkeley (UC Berkeley), Berkeley, CA, USA, in 2017.

He is currently a Research Assistant Professor with the Electrical and Computer Engineering Department, University of Illinois at Urbana-Champaign, Champaign, IL, USA. From 2017 to 2018, he was a Postdoctoral Researcher with the Berkeley Energy and Climate Institute, Berkeley, CA, USA, and a Visiting Lecturer with the Industrial Engineering and Operations Research Department, UC Berkeley. During his Ph.D. studies, he was funded in part by the National Science Foundation Graduate Research Fellowship.