

TESLA: An Extended Study of an Energy-saving Agent that Leverages Schedule Flexibility

Jun-young Kwak · Pradeep Varakantham ·
Rajiv Maheswaran · Yu-Han Chang · Milind
Tambe · Burcin Becerik-Gerber · Wendy Wood

Received: date / Accepted: date

Abstract This paper presents TESLA, an agent for optimizing energy usage in commercial buildings. TESLA's key insight is that adding flexibility to event/meeting schedules can lead to significant energy savings. This paper provides four key contributions: (i) online scheduling algorithms, which are at the heart of TESLA, to solve a stochastic mixed integer linear program (SMILP) for energy-efficient scheduling of incrementally/dynamically arriving meetings and events; (ii) an algorithm to effectively identify key meetings that lead to significant energy savings by adjusting their flexibility; (iii) an extensive analysis on energy savings achieved by TESLA; and (iv) surveys of real users which indicate that TESLA's assumptions of user flexibility hold in practice. TESLA was evaluated on data gathered from over 110,000 meetings held at nine campus buildings during an eight month period in 2011–2012 at the University of Southern California (USC) and the Singapore Management University (SMU). These results and analysis show that, compared to the current systems, TESLA can substantially reduce overall energy consumption.

Keywords Energy · Sustainable Multiagent Systems · Energy-oriented Scheduling · Scheduling Flexibility

Jun-young Kwak · Rajiv Maheswaran · Yu-Han Chang · Milind Tambe
University of Southern California, Los Angeles CA 90089
E-mail: {junyounk,maheswar,ychang,tambe}@usc.edu

Pradeep Varakantham
Singapore Management University, Singapore 178902
E-mail: pradeepv@smu.edu.sg

Burcin Becerik-Gerber · Wendy Wood
University of Southern California, Los Angeles CA 90089
E-mail: {becerik,wendy.wood}@usc.edu

1 Introduction

Reducing energy consumption is an important goal for sustainability. Thus, conserving energy in commercial buildings is important as it is responsible for significant energy consumption. In 2008, commercial buildings in the U.S. consumed 18.5 QBTU¹, representing 46.2% of building energy consumption and 18.4% of U.S. energy consumption [21]. This energy consumption is significantly affected by a large number of meetings or events in those buildings. The main drivers of such energy consumption include HVAC (Heating, Ventilation, and Air Conditioning) systems, lighting and electronic devices. Furthermore, a recent study shows that meeting frequency in commercial buildings is significant and continues to grow [15]. In 2001, U.S. Fortune 500 companies are estimated to have held 11 million formal meetings daily and 3 billion meetings annually.

Energy-oriented scheduling can assist in reducing such energy consumption [4, 36, 46]. Although conventional scheduling techniques compute the optimal schedule for many meetings or events while satisfying their given requirements (i.e., computing a valid schedule) [12, 22, 31, 40], they have not typically considered energy consumption explicitly. More recently, there have been some trials to conserve energy by consolidating meetings in fewer buildings [26, 32]. In particular, Portland State University consolidated night and weekend classes, which were previously scattered across 21 buildings, into five energy efficient buildings. By doing this, they reported that electricity consumption was reduced by 18.5% (78,000 kWh) in the autumn compared to the previous three-year average. Similarly, Michigan State University consolidated classes and events into fewer buildings on campus, and energy reductions in the seven buildings ranged from 2–20%, saving \$16,904. However, these efforts have been decided manually, and no underlying intelligent system was used.

Motivated by this prior work, we describe TESLA (Transformative Energy-saving Schedule-Leveraging Agent), an agent for optimizing energy use in commercial buildings. TESLA is a goal-seeking (to save electric energy), continuously running autonomous agent. TESLA's key insight is that adding *flexibility*, which is a novel concept for capturing user scheduling constraints, to meeting schedules can lead to significant energy savings. Users in a commercial building continuously submit meeting requests to TESLA while indicating flexibility in their meeting preferences. TESLA schedules these meetings in the most energy efficient manner while ensuring user comfort; but in cases where shifting meeting times can lead to significant savings, TESLA interacts with users to request such a shift.

Based on TESLA, this paper provides four key contributions. First, it provides online scheduling algorithms, which are at the heart of TESLA, and presents the sample average approximation (SAA) method [2, 29] to solve a two-stage stochastic mixed integer linear program (SMILP). This SMILP considers the flexibility of people's preferences for energy-efficient scheduling of incrementally/dynamically arriving meetings and events. In this work, flexibility specifically refers to the number of options made available by the user-specified scheduling constraints in terms of

¹ QBTU indicates Quadrillion BTU, which is used as the common unit to explain global energy use. 1 BTU = 0.00029 kWh.

starting time, locations and the deadline before committing to the finalized schedule details. Second, TESLA also includes an algorithm to effectively identify key meetings that could lead to significant energy savings by adjusting their flexibility. Third, this paper provides an extensive analysis of the energy saving results achieved by TESLA. Lastly, surveys of real users are provided indicating that TESLA's savings can be realized in practice by effectively leading people to change their schedule flexibility. To validate our work, we used a public domain simulation testbed [21], fitted it with details of our testbed building, and compared the simulation results against real-world energy usage data. Our results show that, in a validated simulation using our testbed building, TESLA is projected to save about 94,000 kWh of energy (roughly \$18K) annually. Thus, TESLA can potentially offer energy saving benefits to all commercial buildings where meetings affect energy usage.

Although we have focused on evaluating TESLA in commercial buildings, TESLA can be applied to general scheduling domains where schedule flexibility plays a key role for conserving energy. For instance, while scheduling home appliances in residential buildings [4, 28, 36, 44, 46], agents may consider people's preferences and effectively adjust their appliance schedules (e.g., avoid running appliances during peak hours) in order to save energy. Scheduling decentralized appliances in the smart grid can be framed as a decentralized agent-based coordination problem, which can be extended in distributing TESLA [13, 14, 42]. In addition, flexible scheduling could be adopted for manufacturing systems as well [10, 38, 39] as there are many different sets of constraints/preferences while scheduling resources. In particular, specific temporal constraints on some activities such as local release dates (availability of raw material) or due-dates (status review dates) are often flexible in practice, and durations of activities are sometimes controllable, hence are a matter of preference as well.

The rest of the paper is organized as follows: In Section 2, we describe our testbed buildings along with real data from those buildings. In Section 3, we describe the TESLA system and the scheduling algorithms at the heart of it. Section 4 provides evaluations for each of our algorithms using real-world meeting and energy data which indicate that TESLA could potentially provide significant savings in overall energy consumption. Section 5 discusses why TESLA works in detail by providing an extensive analysis on energy savings. In Section 6, surveys of real meeting participants are provided. Section 7 discusses a number of related approaches for handling energy-aware scheduling. We conclude this paper in Section 8.

This paper extends our AAMAS main track paper [19] and features a significant amount of new material. First, Section 3.2.1 now formulates the online scheduling problem as a stochastic mixed integer linear program (SMILP) in order to consider various types of uncertainties as well as people's flexibility. This formulation is more expressive compared to our previous work, which presented a two-step process that attempted to simulate an SMILP. In addition, we include a more general solution technique based on the sample average approximation (SAA) method to solve an SMILP. Thus, we reran all the experiments using the SAA method with the same parameter settings that were used in [19] and report new results in Section 4. Second, Section 3.2.2 discusses the algorithm to identify key meetings not only independently as in [19] but also simultaneously as a group. As we have shown in the evaluation




Fig. 1 The actual research testbed (library) at USC

section, this change has significant potential to improve the overall energy savings. Third, an extensive analysis on energy savings achieved by TESLA is provided in Section 5, which will help readers understand why TESLA works in detail. This analysis was absent in [19]. Fourth, analysis based on the collected meeting request data from SMU (Figure 4) was included in Section 2.2. This analysis was not present in our previous work. Fifth, Section 4 has been significantly enhanced with several new major results. In addition to the new results on SAA that we discussed above, major new additional results include: (i) the scalability and accuracy analysis of the SAA method to solve the SMILP (Section 4.1.2); (ii) energy saving results based on a different prediction heuristic for the predictive non-myopic method (SAA) (Section 4.1.3); (iii) energy saving improvements when simultaneously identifying multiple key meetings (Section 4.1.4); (iv) further energy savings utilizing the cancellation rate of meeting requests (Section 4.1.5); and (v) energy validation results to verify the simulation environment (Section 2.1) and full energy saving results based on real meeting data collected from SMU (Section 4.1.3). Sixth, Section 6 has been strengthened by providing (i) a set of questions used in questionnaire in our human subject experiments (Tables 8 – 10) and (ii) further discussion regarding potential strategic behaviors between agents and human users while focusing on a truthful and fair mechanism at the end of Section 6. There are thus six significant areas of significant improvement over our previous paper.



Fig. 2 The current room reservation system at the testbed building

2 Research Testbed

2.1 Educational Building Testbed


Our system is to be deployed in an educational building. Figure 1 shows the testbed building for TESLA’s deployment and the floor plans of 2nd and basement floors. It is one of main libraries at the University of Southern California and has been designed with a building management system. It hosts a large number of meetings (about 300 unique meetings per regular day) across 35 group study rooms. Each study room has different physical properties including different types and numbers of devices and facilities (e.g., video conferencing equipment, computer, projector, video recorder, office electronic devices, etc.), room size, lighting specification, and maximum capacity (4 – 15 people). This building operates these study rooms 24 hours a day and 7 days a week except on national holidays. The temperature in group study rooms is regulated by the facility managers according to two set ranges for occupied and unoccupied periods of the day. HVAC systems always attempt to reach the pre-set temperature regardless of the presence of people and their preferences in terms of temperature. Lighting and appliance devices are manually controlled by users.

In this building, meetings are requested by users by a centralized online room reservation system (see Figure 2). In the current reservation system, no underlying intelligent system is used; instead, users reactively make a request based on the availability of room and time when they access the system. While users make a request using the system, they are asked about additional information including the number of meeting attendees and special requirements. Reservations can be made up to 7 days in advance.

Given the significant number of meetings per day and the centralized online meeting reservation system, this library testbed provides a good environment to test various energy-oriented scheduling techniques to mitigate energy consumption. TESLA’s goal is to enable users to input flexibility in their scheduling request, to identify key scheduling requests, and to use this information in algorithms that can provide energy-efficient schedules to effectively conserve energy in commercial buildings. To evaluate TESLA, we have built upon a simulation testbed [21] using real building

Table 1 Energy consumption validation (kWh)

Period	Regular semester (Spring/Fall)	Summer break	Average
Actual energy consumption	740.2	289.6	546.7
Simulated energy consumption	721.3	255.1	521.1
Average error (%)	2.6	11.9	4.7

**Fig. 3** Real data analysis (USC)

data and validated with real-world energy data (see Table 1). We specifically compared the energy consumption calculated in the simulation testbed with actual energy meter data from the testbed building (library) at the University of Southern California in 2012. As shown in Table 1, the average difference between actual energy meter data and energy use from the simulation testbed was 4.7%, which strongly supports our claim that the simulation testbed is realistic. This validated simulation environment is used to evaluate TESLA with real meeting data. In addition, we also test TESLA on buildings at the Singapore Management University. SMU has a centralized web-based system that allows users to schedule meetings and events in over 500 conference/meeting rooms across eight buildings. More details regarding the data sets from USC and SMU to test TESLA are provided in the next section.


2.2 Data Analysis

In collaboration with building system managers, we have been collecting data specifying the past usage of group study rooms, which are collected for 8 months (January through August in 2012) at USC. The data for each meeting request includes the time of request, starting time, time duration, specified room, and group size. The data set contains 32,065 unique meetings, and their average meeting time duration is 1.78 hours.

Figure 3(a) shows the actual meeting frequency (y-axis) over time (24 hours, x-axis) of sampled 4 locations at USC (out of 35 rooms) based on the collected meeting

Table 2 Meeting request arrival distribution

Time period	Likelihood (%)
1 day before	55.73
1-2 days before	18.40
2-3 days	8.72
3-4 days	5.52
4-5 days	3.68
5-6 days	3.05
6-7 days	3.35
> 7 days	1.56

**Fig. 4** Real data analysis (SMU)

request data. This figure shows the preferred slots of time and location (e.g., late afternoon (2–5pm) for time & 2nd floor (201A, 202E) compared to the basement for location). Then, the system will be able to predict future situations based on this frequency data while scheduling requests as they arrive. Figure 3(b) shows the probability distribution over total meeting requests per day. The x-axis of the figure indicates the total number of meeting requests per day (ranging from 0 to about 350) and the y-axis shows how likely the system will have the given number of total meeting requests (x-axis) on one day. One can see that the probability of having 50 or fewer meetings is 42.92% and the probability of having 250 or more meetings is 30.04%. These are used to estimate the model of future meetings in our algorithm that will be presented in Section 3.2.

Table 2 shows how early meeting requests were made. In the table, column 2 indicates the percentage of meetings that were requested within the given time period (column 1). For instance, 55.73% of all meeting requests were made within 1 day before the actual meeting day. This analysis would be helpful in understanding how our algorithm could achieve significant energy savings in this domain.

While evaluating TESLA, we also consider another data set from SMU. The data set contains over 80,000 meetings that have been collected for three months (August through October) in 2011 at SMU, which gives us a sense regarding how TESLA




Fig. 5 TESLA architecture: TESLA is a continuously running agent that supports four key features: (i) energy-efficient scheduling; (ii) identification of key meetings; (iii) learning of user preferences; and (iv) communication with users.

will handle energy-oriented scheduling problems in large buildings. Similar to Figure 3, Figure 4(a) shows the actual meeting frequency (y-axis) over time (24 hours, x-axis) of sampled 4 locations at SMU (out of over 500 rooms) based on the collected meeting request data. This figure shows the preferred slots of time and location. Figure 4(b) shows the probability distribution over total meeting requests per day. The x-axis of the figure indicates the total number of meeting requests per day (ranging from 0 to about 1200) and the y-axis shows how likely the system will have the given number of total meeting requests (x-axis) on one day.

3 TESLA

In this section, we describe the overall architecture of TESLA and how to optimally schedule meetings in real-world situations to conserve energy in commercial buildings.

3.1 TESLA Architecture

TESLA is a goal-seeking (to save energy), continuously running autonomous agent. TESLA performs on-line energy-efficient scheduling while considering dynamically arriving inputs from users; these dynamic inputs make the scheduling complex and TESLA needs to learn a predictive model for users' inputs and preferences (see Figure 5). More specifically, TESLA:

- takes inputs (i.e., preferred time, location, the number of meeting attendees, etc.) from different users and their proxy agents at different times (Sections 2 & 3.1)
- autonomously performs on-line energy-efficient scheduling as requests arrive while balancing user comfort (Section 3.2.1)
- autonomously, on own initiative, interacts with different users based on identified problematic key meetings in order to avoid bother cost to users while persuading them to change meeting flexibility (Section 3.2.2)
- bases its non-myopic optimization on learned patterns of meetings (Sections 3.2.1 & 4)

As shown in Figure 5, meeting requests are the information we get from the interface of TESLA via the web interface (or via a proxy agent [34] on an individual user’s hand-held device, in case the users have proxy agents, who have the corresponding users’ preferences and behavior models with a certain level of adjustable autonomy). TESLA focuses on minimizing unnecessary interactions by detecting a small number of key meetings while negotiating with people to adjust their flexibility. TESLA may interact with users’ proxy agents instead of the users themselves.

3.2 TESLA Algorithms

The objective of this work is to come up with energy efficient schedules in commercial buildings with a large number of meetings while considering (i) flexibility in meeting requests over time, location and deadline; and (ii) user preferences with respect to energy and satisfaction. To account for these two constraints, we provide two types of algorithms, which are at the heart of TESLA. First, we provide algorithms that compute a schedule for known and predicted meeting requests which have flexibility in time, location and deadline. Second, based on the schedule obtained, we provide algorithms that detect meeting requests which if modified (to increase flexibility) can result in significant energy savings.

3.2.1 Scheduling algorithms

Before describing our scheduling algorithms, we formally describe the scheduling problem. Let T represent the entire set of time slots available and L represent the set of available locations each day. A schedule request r_i is represented as the tuple: $r_i = \langle a_i, T_i, L_i, \delta_i, d_i, n_i \rangle$, where: a_i is the arrival time of the request, $T_i \subset T$ is the set of preferred time slots for the start of the event and $L_i \subset L$ is a set of preferred locations. d_i is the deadline by which the time and location for the meeting should be notified to the user, δ_i is the duration for the event and finally, n_i is the number of attendees.

The flexibility of the meeting request r_i is a tuple denoted by $\alpha_i: \langle \alpha_i^T, \alpha_i^L, \alpha_i^d \rangle$.²

² Flexibility is already present in the meeting request as its constraints, and α is a measure of such constraints.




Fig. 6 Disjoint sets of R

- α_i^T : time flexibility of meeting i . $\alpha_i^T = \frac{|T_i|-1}{|T|-\delta_i} \times 100$ ($|T| > \delta_i$; i.e., $|T|$ is 24 hours per day).
- α_i^L : location flexibility of meeting i . $\alpha_i^L = \frac{|L_i|-1}{|L|-1} \times 100$ ($|L| > 1$).
- α_i^d : deadline flexibility of meeting i . $\alpha_i^d = \frac{d_i^*-a_i}{d_i^*-a_i} \times 100$, where d_i^* is the latest notification time (e.g., midnight on the meeting day) ($d_i^* > a_i$). $0 \leq \alpha_i^d \leq 100$

For instance, given only one time slot ($|T_i| = 1$), $\alpha_i^T = 0$ and all available time slots ($|T_i| = |T| - \delta_i + 1$), $\alpha_i^T = 100$. Assuming that people give $T_i = 4\text{--}7\text{pm}$ on Monday and their meeting time duration is 2 hours, then $\alpha_i^T = (4-1)/(24-2) \times 100 = 13.64\%$. Likewise, given only one location slot ($|L_i| = 1$), $\alpha_i^L = 0$ and given all available locations ($|L_i| = |L|$), $\alpha_i^L = 100$.

We now define specific disjoint sets of meeting requests, R , that will enable us to characterize different types of scheduling algorithms, where t is the time to schedule a given set of requests R .

- $R^S(t) = \{i : d_i = t \text{ and } a_i \leq t\}$: a set of requests that have to be scheduled at time t
- $R^A(t) = \{i : d_i < t \text{ and } a_i < t\}$: a set of requests that were assigned before time t
- $R^K(t) = \{i : d_i > t \text{ and } a_i \leq t\}$: a set of known future requests, which arrived before time t , but will be scheduled in the future
- $R^U(t) = \{i : d_i > t \text{ and } a_i > t\}$: a set of unknown future requests

As a simple example (shown in Figure 6), let us consider that we have 4 meeting requests (r_1, r_2, r_3 , and r_4), which are supposed to be scheduled on the same day. The current time is t . According to the definition, $R^S(t) = \{r_2\}$, $R^A(t) = \{r_1\}$, $R^K(t) = \{r_3\}$, and $R^U(t) = \{r_4\}$.

Given a set of requests, R , we provide a two-stage stochastic mixed integer linear program (SMILP) to compute a schedule that minimizes the overall energy consumption. Stochastic programming has provided a framework for modeling optimization problems that involve uncertainty [6, 9, 16, 35]. Whereas deterministic optimization problems are formulated with known parameters, real world problems almost invariably include some unknown parameters. In particular, our scheduling problem aims to optimally schedule incrementally/dynamically arriving requests, and thus we should consider uncertainty in terms of future requests, which makes deterministic

optimization techniques inapplicable. To address this challenge, we specifically formulate our scheduling problem as a two-stage stochastic program. Here the decision variables are partitioned into two sets. The first stage variables are decided before the actual realization of the uncertain parameters are known. Afterward, once the random events have exhibited themselves, further decisions can be made by selecting the values of the second stage. The second stage decision variables can be made to minimize penalties that may occur as a result of the first stage decision. This SMILP will be run every time a new meeting request arrives (or after a batch of meeting requests arrive in close succession).

The notation that will be employed in the SMILP is as follows:

- $x_{l,t}^i$ is the first stage binary variable that is set to 1 if meeting request r_i is scheduled in location l starting at time t .
- $E_{l,t}^i$ is a constant that is computed for a meeting request r_i if it is scheduled in location l at time t using the HVAC energy consumption equations.
- C is a constant that indicates the reduction in energy consumption because of scheduling a meeting in the previous time slot. Although we assumed that C is a constant for simplicity in this work, it depends on different factors of previous meetings in practice.
- $e_{l,t}^i$ is a continuous variable that corresponds to the energy consumed because of scheduling meeting i in location l at time t . The value of this variable is affected based on whether there is a meeting scheduled in the previous time slot ($t-1$), i.e., the reduction that would occur at location l at time t if a meeting was scheduled at location l at time $t-1$.³ $e_{l,t}^i = x_{l,t}^i \cdot E_{l,t}^i - \sum_{i' \in R \setminus \{i\}} x_{l,t-1}^{i'} \cdot C$.
- $S_{l,t}^i$ is a value that indicates the satisfaction level obtained with users in meeting request r_i for scheduling the meeting in location l at time t . B is a threshold on the satisfaction level required by users.
- M is an arbitrarily large positive constant.
- $Q(x, \xi)$ is the value function of future energy consumption, where ξ represents uncertainty over the second stage problem (i.e., future meeting situations in our problem). ξ determines a vector of parameters, (w, q) .
- $w_{l,t}^j$ is the second stage binary variable that is set to 1 if meeting request r_j in a future meeting request set is scheduled in location l starting at time t .
- $q_{l,t}^j$ is a continuous second stage variable that corresponds to the future energy consumed because of scheduling meeting j in location l at time t .

We first provide the SMILP and a detailed explanation of the constraints.

³ $e_{l,t}^i$ gets affected by a meeting in the previous time slot in the same location. This is because adjacent meetings affect the indoor temperature, which makes HVACs operate differently to maintain the desired temperature level.

$$\min e + \mathbb{E}[Q(x, \xi)] \quad (1)$$

{Choose the optimal first stage variables that minimizes the sum of first stage costs and the expected value of the second stage}

s.t.

$$e \geq \sum_{i \in R \setminus R^U} \sum_{t \in T} \sum_{l \in L} e_{l,t}^i, \quad (2)$$

{Computing the first stage cost e }

$$e_{l,t}^i = x_{l,t}^i \cdot E_{l,t}^i - \sum_{i' \in R \setminus R^U \setminus \{i\}} x_{l,t-1}^{i'} \cdot C, \quad \forall i \in R \setminus R^U, l \in L, t \in T \quad (3)$$

{Computing energy consumption while considering the back-to-back meeting effect}

$$e_{l,t}^i \geq 0, \quad \forall i \in R \setminus R^U, l \in L, t \in T \quad (4)$$

$$\sum_{t \in T} \sum_{l \in L} x_{l,t}^i \cdot S_{l,t}^i \geq B, \quad \forall i \in R \setminus R^U \quad (5)$$

{Checking if the computed schedule maintains the given comfort level B }

$$\sum_{i \in R \setminus R^U} x_{l,t}^i \leq 1, \quad \forall l \in L, t \in T \quad (6)$$

$$\sum_{i' \in R \setminus R^U \setminus \{i\}} \sum_{t'=t}^{t+\delta_i-1} x_{l,t'}^{i'} \leq M(1 - x_{l,t}^i), \quad \forall l \in L, i \in R \setminus R^U, t \in T \quad (7)$$

{Checking the allocation restrictions that for each assignment slot, only one meeting can be scheduled considering the given time duration of meeting}

$$x_{l,t}^i \in \{0, 1\}, \quad \forall i \in R \setminus R^U, l \in L, t \in T \quad (8)$$

{The first stage binary variable}

$$Q(x, \xi) \geq \sum_{j \in R^U} \sum_{l \in L} \sum_{t \in T} q_{l,t}^j, \quad (9)$$

{Computing the second stage cost Q }

$$q_{l,t}^j = w_{l,t}^j \cdot E_{l,t}^j - \sum_{i \in R \setminus R^U} x_{l,t-1}^i \cdot C - \sum_{i \in R \setminus R^U} x_{l,t+1}^i \cdot C - \sum_{j' \in R^U \setminus \{j\}} w_{l,t-1}^{j'} \cdot C, \quad (10)$$

{Computing energy consumption while considering the back-to-back meeting effect caused by the first and second stage variables}

$$q_{l,t}^j \geq 0, \quad \forall j \in R^U, l \in L, t \in T \quad (11)$$

$$\sum_{j \in R^U} w_{l,t}^j \leq 1, \quad \forall l \in L, t \in T \quad (12)$$

$$\sum_{j \in R^U} \sum_{t'=t}^{t+\delta_j-1} w_{l,t'}^j \leq M(1 - x_{l,t}^i), \quad \forall l \in L, i \in R \setminus R^U, t \in T \quad (13)$$

{Checking the allocation restrictions against the first stage assignment slots}

$$\sum_{j' \in R^U \setminus \{j\}} \sum_{t'=t}^{t+\delta_j-1} w_{l,t'}^{j'} \leq M(1 - w_{l,t}^j), \quad \forall l \in L, j \in R^U, t \in T \quad (14)$$

{Checking the allocation restrictions against the second stage assignment slots}

$$w_{l,t}^j \in \{0, 1\}, \quad \forall j \in R^U, l \in L, t \in T \quad (15)$$

{The second stage binary variable}

The objective of the SMILP above is to choose the optimal first stage variables (i.e., the optimal assignment of meeting requests to locations and time slots that is characterized by the solution, $x_{l,t}^{i*}$). The optimal first stage variable, x^* , is selected in a way that the sum of first stage costs e (i.e., the energy consumption when the current meeting request is scheduled) and the expected value of the second stage or recourse costs $\mathbb{E}[Q(x, \xi)]$ (i.e., the expected energy consumption that will be realized by future meeting requests) is minimized. In this formulation, at the first stage we have to make a decision before the realization of the uncertain data ξ , which is viewed as a random vector that determines future meeting requests, is known. At the second stage, after a realization of ξ becomes available, we optimize our behavior by solving an appropriate optimization problem.

Constraints (2) – (8) are a set of enforcement for deciding first stage variables, and constraints (9) – (15) enforce conditions for second stage variables. More specifically, constraint (3) is for computing energy consumption considering the back-to-back meeting effect. In particular, we subtract from the energy consumed by this meeting indexed by i at time t , the impact due to meetings (indexed by i'), that were scheduled at the prior time slot $t - 1$. Constraint (5) is for checking if the computed schedule maintains the given comfort level B . Constraints (6) and (7) are the allocation restrictions that for each assignment slot, only one meeting can be scheduled considering the given time duration of meeting. In particular, M in constraint (7) is an arbitrarily large positive constant to enforce only one meeting is scheduled at a location during the duration of the meeting. If meeting i is assigned to location l and time t ($x_{l,t}^i = 1$), then any other meeting requests cannot be assigned to the same slot. If $x_{l,t}^i = 0$, the constraint does not block any other meeting requests from being assigned to that slot as the right-hand side of the equation is not bounded due to an arbitrarily large constant of M . Constraint (9) is to compute the optimal value of the second stage problem while satisfying constraints (10) – (15) which are similar to constraints (3) – (8). Specifically, constraint (10) is for computing the energy reduction that would occur if there are any consecutive meetings among the requests in R^U (i.e., check with w) and if any future meetings have this back-to-back effect with either already assigned meetings or ones that have to be scheduled in $R \setminus R^U$ (i.e., check with x).

We now describe the sample average approximation (SAA) method [2,29] to solve the given SMILP. The main idea of the SAA approach to solve stochastic programs is as follows. A sample ξ^1, \dots, ξ^N realizations of the random vector ξ is generated, and consequently the expected value function $\mathbb{E}[Q(x, \xi)]$ in the stochastic program (1) is approximated by the weighted average function $\sum_{n=1}^N p_n^U Q(x, \xi^n)$, where p_n^U is the likelihood that ξ^n is realized. Recall that ξ is the random vector that determines future meeting requests in our formulation (i.e., each realization ξ^n has a different number of future meeting requests and corresponding request tuples). More specifically, we have a probability distribution p^T over the possible range of total meeting requests per day (shown in Figures 3(b) & 4(b)). Then, the likelihood that k more meetings will arrive on the same day assuming we currently have s meetings so far is equivalent to the likelihood that ξ^n is realized with k unknown future requests: $p_n^U(k) = p^T(s + k)$. For those k future meeting requests in R_n^U , we generate random

request tuples (specifically, T_i & L_i) based on the actual distribution over the assignment spots as shown in Figures 3(a) & 4(a). Then, for a sample n ($1 \leq n \leq N$), the original SMILP is reformulated as follows:

$$\min e + \sum_{n=1}^N p_n^U Q(x, \xi^n) \quad (16)$$

{Using SAA, the expected value of the second stage cost is approximated by the weighted average function. Then, we still choose the optimal first stage variable that minimizes the sum of the first and second stage costs}

s.t.

Constraints (2) – (8),

$$Q(x, \xi^n) \geq \sum_{j \in R_n^U} \sum_{l \in L} \sum_{t \in T} q_{j,l,t}^n, \quad (17)$$

$$q_{j,l,t}^n = w_{j,l,t}^n \cdot E_{j,l,t} - \sum_{i \in R \setminus R^U} x_{i,l,t-1} \cdot C - \sum_{i \in R \setminus R^U} x_{i,l,t+1} \cdot C - \sum_{j' \in R_n^U \setminus \{j\}} w_{j',l,t-1}^n \cdot C, \quad (18)$$

$$q_{j,l,t}^n \geq 0, \quad \forall j \in R_n^U, l \in L, t \in T \quad (19)$$

$$\sum_{j \in R_n^U} w_{j,l,t}^n \leq 1, \quad \forall l \in L, t \in T \quad (20)$$

$$\sum_{j \in R_n^U} \sum_{t'=t}^{t+\delta_i-1} w_{j,l,t'}^n \leq M(1 - x_{i,l,t}), \quad \forall l \in L, i \in R \setminus R^U, t \in T \quad (21)$$

$$\sum_{j' \in R_n^U \setminus \{j\}} \sum_{t'=t}^{t+\delta_j-1} w_{j',l,t'}^n \leq M(1 - w_{j,l,t}^n), \quad \forall l \in L, j \in R_n^U, t \in T \quad (22)$$

$$w_{j,l,t}^n \in \{0, 1\}, \quad \forall j \in R_n^U, l \in L, t \in T \quad (23)$$

$$\sum_{n=1}^N p_n^U = 1 \quad (24)$$

{ p_n^U is the likelihood that ξ^n is realized, where ξ is a random variable that determines future meeting requests U }

The obtained sample average approximation (16) of the stochastic program is then solved using a standard branch and bound algorithm such as those implemented in commercial integer programming solvers such as CPLEX.

As benchmark algorithms for comparison purposes, we provide two optimization heuristics: myopic and full-knowledge. We have the myopic optimization algorithm, which obtains a schedule by considering the following request set: $R = (R^A(t) \cup R^S(t) \cup R^K(t))$. A schedule and energy consumption are obtained without

accounting for future unknown meetings. Thus, the myopic heuristic only considers the first stage decision variables in our SMILP. In the full-knowledge method, we compute the final schedule while assuming that the entire set of meeting requests R is given, which is ideal. Thus, for the full-knowledge method, we have one actual realization with probability 1.0 for computing the second stage costs in the SMILP. The performance comparison results will be provided in Section 4.

3.2.2 Identifying key meetings

TESLA computes the optimal schedule considering the given flexibility (or scheduling constraints) of meetings. It can obtain more energy-efficient schedules by increasing flexibility (i.e., relaxing those constraints). We now provide an algorithm that finds meeting requests, which if made more flexible will reduce energy consumption significantly.

Algorithm 1 IDENTIFYKEYMEETINGS (\mathbf{R})

```

1:  $\mathbf{U} \leftarrow \emptyset$ 
2: {Initialize a set of key meetings}
3:
4: for all  $\mathbf{I} \subset 2^R$  do
5:   { $R$  is a set of requests.}
6:   if ISAVINGCANDIDATE ( $\mathbf{I}$ ) then
7:      $\mathbf{U} \leftarrow \mathbf{U} \cup \mathbf{I}$ 
8:
9: return  $\mathbf{U}$ 

```

Algorithm 1 describes the overall flow of the algorithm. We first initialize a set that will contain key meetings identified by our algorithm (line 1). For each subset of the power set of meeting requests R , we then examine whether or not the current meeting set I is a key meeting set by relying on Algorithm 2 (line 6).

Algorithm 2 recursively determines if the given meeting set I is a candidate set that gives significant potential energy savings. The meeting set I is detected as a key meeting set only if the expected energy savings of meeting requests in I are monotonically increasing and show higher energy improvements than the given threshold value (τ ; a certain level of additional energy savings that we desire to achieve with the selected key meetings) by relaxing their flexibility. To handle this, we first compute the expected energy savings of the meeting set I when its flexibility level is changed from the initial level α_I to the desired level α'_I assuming the other meetings' flexibility levels are fixed (line 1). The expected energy saving value of meeting set I , $V_I = (E_{\alpha_I} - E_{\alpha'_I})/E_{\alpha_I}$ ($0 \leq V_I \leq 1$), where E_{α_I} is the current total energy consumption with the given level of flexibility α_I , and $E_{\alpha'_I}$ is the reduced total energy consumption if the meeting set I 's flexibility is changed to one of k possible options, $\alpha'_{I,k}$, while others keep their given flexibility levels. In this work, we consider a heuristic for setting the threshold value to investigate whether or not the current meeting set I is an energy saving candidate set: a fixed single threshold value τ (line 5; e.g., 0.4 as a universal threshold).

Algorithm 2 ISSAVINGCANDIDATE (I)

```

1:  $V_I \leftarrow \text{CALEXPENERGYSAVINGS}(\alpha_I, \{\alpha'_{I,1}, \dots, \alpha'_{I,k}\})$ 
2:  $\{\alpha_I$  is an initially given flexibility of meetings in  $I$ , and  $\alpha'_{I,k}$  is one of the desired flexibility options
   for meetings in  $I$ . CALEXPENERGYSAVINGS computes energy gains,  $V_I$ , by relaxing flexibility of
   meeting requests in  $I$ . $\}$ 
3:
4: if  $|I| = 1$  then
5:   if  $V_I > \tau$  then
6:     {If the computed energy gains  $V_I$  is higher than a given threshold value  $\tau$ , it is considered as a
       key meeting.}
7:     return TRUE
8:   else
9:     return FALSE
10: else if  $|I| > 1$  then
11:   {Recursively call ISSAVINGCANDIDATE with possible subsets}
12:   for all  $i \in I$  do
13:      $I' \leftarrow I \setminus \{i\}$ 
14:      $V_{I'} \leftarrow \text{CALEXPENERGYSAVINGS}(\alpha_{I'}, \{\alpha'_{I',1}, \dots, \alpha'_{I',k}\})$ 
15:     if  $V_I - V_{I'} > 0$  then
16:       {Only if the energy savings are monotonically increasing by adding a meeting request  $i$  (or
         monotonically decreasing by excluding a meeting request  $i$ ), proceed}
17:       return ISSAVINGCANDIDATE ( $I'$ )

```

4 Empirical Validation

We evaluate the performance of TESLA and experimentally show that it can conserve energy by providing more energy-efficient schedules in commercial buildings. At the end of this section, we provide actual survey results that we have conducted on schedule flexibilities of real users. The experiments were run on Intel Core2 Duo 2.53GHz CPU with 8GB main memory. We solved our MILP formulations using CPLEX version 12.1. All techniques were evaluated for 100 independent trials and we report the average values. Energy consumption was computed using the simulator described earlier in Section 2.1.

4.1 Simulation Results

In this section, we provide the simulation results (i) to verify if flexibility *really* helps TESLA compute energy-efficient schedules; (ii) to extensively evaluate the overall performance of the SAA method while varying the sample size and flexibility; and (iii) to measure energy saving benefits by identifying key meetings and by considering the cancellation rate.

4.1.1 Does flexibility help?

As an important first step in deploying TESLA, we first verified if the agent could save more energy with more flexibility while scheduling given meeting and event requests. To that end, we compared the energy consumption of three different approaches using the real-world meeting data mentioned in Section 2.2: (i) the current




Fig. 7 Energy savings: Actual - the amount of energy consumed in simulation based on the past schedules obtained from the current manual reservation system; Random - energy consumption while randomly perturbing the starting time and location of meeting requests from the same past schedules while keeping meeting time duration; Optimal - Energy consumption measured in simulation based on optimal schedules computed from an SMILP with the fully known meeting request set and full flexibility

benchmark approach in use at the testbed building; (ii) a random method that randomly assigns time and location for meetings; and (iii) the optimal method using the full-knowledge optimization technique described in Section 3.2.

Figure 7 shows the average daily energy consumption in kWh computed based on schedules from the three algorithms above. In the figure, the consumption is the amount of energy consumed based on the past schedules obtained from the current manual reservation system, which shows a very similar performance to the random approach. The optimal method assuming the full amount of flexibility (i.e., 24 hours for α^T , 35 rooms for α^L and delay the deadline before which the final schedule should be informed for α^d) achieved statistically significant energy savings of 50.05% compared to the current energy consumption at the testbed site. These savings are practically significant, and also statistically significant (paired-sample t-test; $p < 0.01$). These savings are equivalent to annual savings of about \$18,600 considering an energy rate of \$0.193/kWh [41] and CO_2 emissions from the energy use of 5.5 homes for one year. Thus, flexibility can help save energy.

4.1.2 Online scheduling method with flexibility: Determining the sample size in the TESLA SMILP

In this section, we first investigated the runtime and solution qualities for solving the SMILP while varying the number of samples (see Figure 8). Figure 8(a) shows the results of the runtime analysis in seconds (y-axis) for sample sizes $N = 10$ to 100 (x-axis). As shown in the figure, the runtime increases in an exponential fashion as the sample size N increases. However, Figure 8(b) shows that its solution quality also increases (y-axis) (i.e., the estimated optimality gap decreases) as the number of samples N increases. For evaluating the generated solution for each of sample size




Fig. 8 Scalability and accuracy while varying the number of samples (N)




Fig. 9 Energy savings while varying flexibility (USC)

N , we generated M independent samples (i.e., replications) of the uncertain parameters, and evaluated the obtained solution in each $m \in M$ replication. In this work, we specifically used 1,000 independent replications for measuring the estimated optimality. Comparing the full-knowledge schedules based on actual realization of each of the 1000 samples with the schedule from the SMILP gives us the percentage error. Based on this result, throughout the paper, we set $N = 50$ to solve the SAA problem. This sample size has a reasonable runtime without a significant compromise in solution quality.

4.1.3 Performance of online scheduling method with flexibility

We next compared solution qualities of the three scheduling algorithms in TESLA presented in Section 3.2.1. Figure 9 shows that how much each algorithm saves when

Table 3 Performance comparison between SAA and myopic

	Max	Min	Average
Optimality difference	57.89%	0.50%	12.73%

compared to the optimal value (i.e., full-knowledge optimization assuming the full flexibility) while varying the time and location flexibility level (assuming 0% deadline flexibility). *The flexibility in our model represents a 3-dimensional space (time, location and deadline), which we have thoroughly explored.* We show results exploring deadline flexibility later.

The optimality percentage on the y-axis of Figure 9 is computed as follows: $(E_a - E_c)/(E_a - E_o)$. Here E_a is the actual energy consumption without any flexibility, E_o is the optimal energy consumption, and E_c is the computed energy consumption using three different scheduling algorithms that we compare using the real meeting data.

Figure 9 shows the average optimality in percentage of each algorithm (M: myopic, P: predictive non-myopic (SAA) and F: full-knowledge) while varying the location flexibility (α^L ; x-axis) and time flexibility (α^T ; each graph assumed the different amount of α^T as indicated in the legend). In the figure, for each pair of flexibility values (α^T, α^L), we report the average optimality in percentage (i.e., 100% indicates the optimal value, and 0% means that there was no improvement from the actual energy consumption). For instance, when flexibility $(\alpha^T, \alpha^L) = (31.5\%, 58.8\%)$, the myopic method achieved an optimality of 50.8%. In the figure, higher values indicate better performance.

As shown in Figure 9, as users provide more flexibility, TESLA can compute schedules with less energy consumption. The gain in optimality from myopic to predictive non-myopic (SAA) is because the latter can leverage user flexibility to put a meeting in a suboptimal spot at the meeting request time to account for future meetings, yielding better results at the actual day of meetings. For example, a flexible meeting request can be moved away from a known popular time-location spot. We conclude that (i) the predictive non-myopic (SAA) method is superior to the myopic method. Table 3 shows the average performance comparison results between the predictive non-myopic (SAA) method and the myopic technique. As shown in the table, the maximum and average optimality differences between the two methods (i.e., optimality of the SAA - optimality of the myopic) are 57.89% and 12.73%, respectively, which are significant. In addition, for 12.50% of cases, the predictive non-myopic (SAA) optimization showed over 20% higher optimality than the myopic method; (ii) the predictive non-myopic (SAA) method performs almost as well as the full-

Table 4 % of optimal energy savings: varying α^T , α^L , and p_f (USC)

		Alg.		Location flexibility (α^L)				
				p_f	23.5	47.1	70.6	94.1
T. flex. (α^T)	0	M	1.0	6.6	6.7	17.8	23.3	
			0.8	5.6	6.0	14.5	21.2	
			0.5	4.9	4.9	13.8	18.2	
			0.2	3.3	3.8	8.4	12.0	
		P	1.0	9.7	9.8	22.7	24.8	
			0.8	8.6	9.3	20.9	23.2	
			0.5	6.4	6.9	15.6	18.6	
			0.2	4.2	4.9	9.8	12.9	
		F	1.0	9.9	10.1	23.6	25.8	
			0.8	8.3	8.6	20.7	24.0	
			0.5	6.7	6.9	16.9	19.1	
			0.2	4.9	5.1	11.3	13.6	
	31.5	M	1.0	46.3	46.5	55.8	61.4	
			1.0	48.1	48.5	62.1	62.7	
		F	1.0	49.0	49.2	63.0	63.1	
			0.8	41.9	43.3	55.5	57.6	
			0.5	29.9	30.7	43.9	44.5	
			0.2	16.1	16.7	26.9	27.2	
		67.5	M	1.0	81.8	82.5	89.6	96.0
				1.0	84.4	86.3	95.4	96.8
	F		1.0	86.3	86.8	96.0	97.5	
			0.8	73.3	73.5	87.9	91.3	
			0.5	53.7	54.4	65.0	67.8	
			0.2	29.4	30.6	38.2	41.4	

(M: myopic, P: predictive non-myopic (SAA), F: full-knowledge)

knowledge optimization (about 98%)⁴; and (iii) full flexibility is not required to start accruing benefits of flexibility.

In the real-world, it is hard to imagine that all people will simply comply and change their flexibility to achieve such optimality. Thus, we provide one additional result shown in Table 4 which varies the percentage of meetings that will have flexibility (p_f). We show α^T along the rows and α^L along the columns. In particular, the value of row 10 and column 5 (highlighted in the table) shows the optimality achieved by the predictive method assuming that 20% of meetings (randomly selected) have $(\alpha^T, \alpha^L) = (0\%, 23.5\%)$ flexibility and the remaining 80% have no flexibility. Our main conclusions are: (i) if we increase p_f , we are able to achieve a higher optimality; and (ii) flexibility in a small number of meetings can lead to significant energy reduction. This motivates considering more intelligent identification of key meetings to change their flexibility (described in the next section).


We also compared the performance of the three algorithms while varying the deadline flexibility, α^d . In Table 5, columns indicate different amounts of deadline

⁴ The average performance of the predictive non-myopic (SAA) optimization depends on the prediction method of future requests. We, thus, additionally tested a more sophisticated prediction method considering the time factor that is one of key features determining the overall trend of requests (i.e., when the meeting requests arrive at the system to be scheduled; e.g., regular semester vs. summer/ winter break). With this additional consideration, the predictive non-myopic (SAA) method improved the overall performance of the predictive method by 1.1%.

Table 5 Percentage of optimal energy savings: varying α^d (USC)

Alg.	α^d	0.0	22.2	44.4	66.7	88.9
M		82.5	83.4	84.0	84.2	84.2
P		86.3	86.4	86.7	86.7	86.8
F		86.8	86.8	86.8	86.8	86.8

(M: myopic, P: predictive non-myopic (SAA), F: full-knowledge)

**Fig. 10** Energy savings while varying flexibility (SMU)**Table 6** Percentage of optimal energy savings: varying α^d (SMU)

Alg.	α^d	0.0	22.2	44.4	66.7	88.9
M		85.30	87.22	89.02	89.41	90.06
P		93.01	93.05	94.56	94.87	95.14
F		95.21	95.21	95.21	95.21	95.21

(M: myopic, P: predictive non-myopic (SAA), F: full-knowledge)

flexibility and values are the optimality of each algorithm assuming a fixed time and location flexibility $(\alpha^T, \alpha^L) = (67.5\%, 47.1\%)$. As we increase the deadline flexibility, both myopic and predictive non-myopic (SAA) methods converge to the full-knowledge optimization result. This is because as the deadline flexibility increases, we can delay scheduling until we have more information. In this particular case of α^T and α^L , we do not necessarily see significant benefits by providing more deadline flexibility since the myopic and predictive non-myopic (SAA) methods already achieved fairly high optimality compared to the full-knowledge method. While the optimality percentage changes are small, given the vast amount of energy consumed by large-scale facilities, these reductions can lead to significant energy savings. We are investigating conditions where our algorithms get more benefits by deadline flexibility.

Table 7 Energy improvement of identified key meetings (%)

$\alpha' \backslash \alpha$	(0,23.5)	(0,47.1)	(0,70.6)	(31.5,23.5)	(31.5,47.1)
(0,23.5)	-	-	-	-	-
(0,47.1)	16.08	-	-	-	-
(0,70.6)	30.08	29.17	-	-	-
(31.5,23.5)	32.05	-	-	-	-
(31.5,47.1)	46.18	36.27	-	29.17	-
(31.5,70.6)	46.52	38.33	34.36	31.07	26.08

The same types of analysis are performed with another data set from SMU and results are presented in Figure 10. The figure shows the average optimality in percentage of each algorithm (M: myopic, P: predictive non-myopic (SAA) and F: full-knowledge) on the y-axis while varying the time flexibility (α^T ; each graph assumed the different amount of α^T as indicated in the legend) and location flexibility (α^L ; x-axis). We assume the deadline flexibility (α^d) of 0%. Similar to earlier results, the predictive method achieved about 97% optimality compared to the full-knowledge optimization and showed higher value than the myopic approach. We also compared the performance of the three algorithms while varying the deadline flexibility. In Table 6, values are the optimality of each algorithm assuming a fixed time and location flexibility, (31.5%, 47.1%). Here we see more pronounced energy savings at SMU as α^d increases compared to the USC results.

4.1.4 Performance of identifying key meetings

We evaluated the performance of the algorithm to identify key meetings for energy reduction. In our tests, we selected 10 meetings individually using the algorithm presented in Section 3.2.2 and calculated the average energy savings if those selected meetings changed their flexibility.

Table 7 shows the average energy savings as described for various flexibility transitions. Columns indicate the initial level of flexibility ($\alpha = (\alpha^T, \alpha^L)$) and rows show the requested level of flexibility ($\alpha' = (\alpha'^T, \alpha'^L)$). For instance, the value in row 4 and column 3 (highlighted in the table) indicates a 29.17% average energy savings improvement if flexibility of 10 key meetings are changed from (0%, 47.1%) to (0%, 70.6%). An important interpretation of that results is that changing the flexibility of key meetings, when those ones are from an appropriately chosen set, contributed to significant energy savings. We also tested how much we can save energy if we choose key meetings simultaneously rather than independently. Assuming the current flexibility is (0%, 23.5%) (column 2 in Table 7), if we choose 10 key meetings at the same time using the same algorithm presented in Section 3.2.2, the average energy savings were improved by 10.3% (i.e., 44.48% of energy saving improvements on average). In the future, we will investigate another heuristic to set a feasible threshold value based on a learned profile of user likelihood of changing meeting flexibility.




Fig. 11 Average energy improvement while considering the cancellation rate of meeting requests

4.1.5 Considering the cancellation rate

According to the real meeting data collected for eight months (January through August in 2012) at USC, about 10.12% (3,245 out of 32,065) of the total meeting requests were canceled, which gives us another insight to achieve further energy savings by utilizing this feature. To incorporate this feature into our SMILP formulation⁵, we change constraint (7) as follows:

$$Pr(\sum_{i' \in R \setminus R^U \setminus \{i\}} \sum_{t'=t}^{t+\delta_i-1} x_{i',l,t'} \leq M(1 - x_{i,l,t})) \geq 1 - \alpha_c$$

The constraint above is given in the form of the chance constrained programming that relaxes the allocation restrictions (i.e., with a probability of α_c , the given allocation restrictions can be violated). In this work, we tested how much additional energy savings can be achieved by allowing the system to overbook meeting rooms that are taken by meeting requests that may be canceled, which is systematically controlled by the cancellation rate (α_c) in the stochastic program. If any schedule conflicts occur by TESLA, TESLA greedily finds the currently available best slots in terms of energy savings for resolving conflict in meetings.

A result is provided in Figure 11. The y-axis in the figure indicates the average energy saving improvements in percentage while varying the cancellation rate (α_c) on the x-axis. These average values were measured over 100 independent trials. As shown in the figure, as we set a higher α_c , the overall average energy savings increase. In particular, with 10.12% cancellation rate that was obtained from the real-world data, the expected energy saving improvement was about 14.78%, which is fairly significant.




Fig. 12 Energy savings by TESLA: the percentage of energy savings per each energy consumer and factor




Fig. 13 Energy saving analysis: room size

5 Analysis: Savings due to TESLA

There are three major components that affect energy consumption in commercial buildings: HVACs (accounting for 35% of the entire energy consumption in commercial buildings), lighting (27%), and electronic devices (about 10%) [18]. TESLA focuses on these three energy consumers to save energy by computing energy-efficient schedules that exploit key factors that affect energy consumption of each building component. Figure 12 shows the percentage of energy savings per each energy consumer and factor in TESLA assuming an actually measured time and location flexibility $(\alpha^T, \alpha^L) = (25.34\%, 16.05\%)$ from surveys of real users. For instance, as shown in the figure, 47.4% of energy savings by TESLA is achieved through more energy-efficient operations of HVACs. More specifically, TESLA shifts meetings to suitable smaller offices or non-peak time and packs meetings together, and those strategies result in a significant energy reduction for HVACs.

⁵ Note that canceled meetings were not considered while scheduling meetings in the earlier results.




Fig. 14 Energy savings only by HVACs (Non-peak Time)

5.1 HVACs

Key assumptions The following assumptions are made in TESLA:

- HVACs are centrally regulated by the university facility management team to satisfy two pre-defined temperature ranges: occupied time zone (8am to 6pm: 70–75F) and unoccupied time zone (rest of the hours: 60–80F).
- While optimizing schedules, the threshold of people’s comfort level was set to 50%, which is a configurable parameter.

Factors impacting HVAC energy As shown in Figure 12, given the above assumptions, HVACs accounted for 47.4% of the overall energy savings. Numbers in the parentheses below indicate the amount of energy savings by each of the following three factors:

- Room Size: TESLA focuses on assigning meetings to smaller spaces while considering the number of meeting attendees, since a larger room requires more energy than a smaller room when occupied for the same amount of time (38.3%). Figure 13 shows the actual and optimal usage density and the physical size (y-axis) of 35 different rooms (x-axis) in the testbed building at USC. As shown in the figure, TESLA generates the schedule that uses 18.16% less space compared to the actual schedule, which clearly proves that TESLA provides more energy-efficient schedules by assigning meetings to smaller spaces.
- Non-peak Time: TESLA avoids the peak time in terms of energy and popularity considering the given constraints/flexibility. Since an unoccupied time zone requires less energy than occupied time zone when the same room is occupied for the same amount of time, TESLA focuses on assigning meetings under an unoccupied time zone as much as possible (29.5%). However, since an unoccupied time zone has a wider regulated temperature range, this optimization may cause a drop in the average comfort level of people. While this flexibility of holding the

meeting at non-peak time is assumed to be part of the meeting request, this drop in comfort level is worth further investigation. The first point to note is that the amount of energy savings achieved by the non-peak time factor itself is less significant (i.e., 13.93%) compared to other factors. Thus, in Figure 14, we provide a result that shows how the non-peak time factor affects the overall energy savings (y-axis) while varying the unoccupied time zone temperature (x-axis). As shown in the figure, as we reduce a temperature range for the unoccupied time zone, the amount of energy savings by the non-peak time factor decreases, but TESLA can still achieve meaningful energy savings while satisfying the given comfort level constraint. Furthermore, TESLA provides a flexible architecture that allows people to configure the temperature value accordingly under different situations.

- Packing Meetings: TESLA focused on packing meetings together in terms of the time interval between meetings in the same room. When a meeting ends, the room is conditioned to a pre-defined environment. This built-up thermal momentum can benefit later meetings scheduled in the same room in close proximity by reducing the number of changes of HVAC operations, which saves much more energy (32.2%).

5.2 Lighting

Key assumptions The following assumptions are made in TESLA:

- The standard nominal values were used for the lighting configuration in spaces.
- When the room was occupied, the full (100%) lighting level was considered.
- When the room was unoccupied, 0% lighting level was considered.

Factors impacting lighting energy As shown in Figure 12, given the above assumptions, the lighting sources accounted for 37.5% of the overall energy savings. The entire energy savings are caused by different room size; specifically, TESLA focuses on assigning meetings to smaller spaces while considering the number of meeting attendees, since a larger room requires more energy than a smaller room when occupied for the same amount of time (see Figure 13).

5.3 Electronics

Key assumptions The following assumptions are made in TESLA:

- Assumed average number of devices in each room was considered to calculate the correct energy consumption.⁶
- When the room was occupied, 80% of the devices were used.
- When the room was unoccupied, 0% of the devices were used.

⁶ While evaluating TESLA, we considered the assumed average number of electronic devices including the actual number of devices existing in each room as well as the average number of devices that people bring with them.

Q1. For each day of the week, please specify your **available** range of starting times for the meeting.

Mon: 8am-3pm



...

Sun: 7am-9am



Q2. What **possible** locations would you consider for your meeting?

201A	201B	201C	201E	201F	202G
202A	B-3D	B-3E	P-3B	BL-3Z	B-3X

Fig. 15 Screenshot of online survey: people were asked to indicate their meeting requests and flexibility.

Factors impacting electronics energy As shown in Figure 12, given the above assumptions, the electronics accounted for 15.1% of the overall energy savings. The entire energy savings are caused by different room size; specifically, TESLA focuses on assigning meetings to smaller spaces while considering the number of meeting attendees, since a larger room has more devices in the testbed building, and thus it requires more energy than a smaller room when occupied for the same amount of time (see Figure 13).

6 Human Subject Experiments

The goal of human subject experiments is to support the results provided in the previous section by answering several questions: (i) are people flexible in real situations?; (ii) how flexible are people in modifying their requests?; (iii) will people in the identified key meetings actually agree to change their flexibility to contribute energy savings?; and (iv) what would be an effective way for an agent to persuade people? To answer these, we measure the amount of reported flexibility change while varying feedback about the energy usage.

We conducted two surveys on a pilot sample of participants (students on campus): (i) an online survey to understand flexibility of those who are using the testbed building; and (ii) a survey to measure flexibility change due to messaging.

6.1 Survey for initial flexibility


We conducted an online survey to understand the flexibility of meeting attendees (shown in Figure 15). The procedure to conduct this survey is as follows: we recruited 32 students who have used the meeting reservation system at the tested building and their facilities. They filled out our survey, indicating meeting requests and flexibility.

Table 8 Basic Profile Questionnaire

Question	Answer (Scale)
Q ₁ . Gender?	Male / Female
Q ₂ . Position at USC?	Undergraduate / Graduate / Staff / Faculty
Q ₃ . Age?	20 or under / 21–25 / 26–30 / 31–35 / 36–40 / 41 or above
Q ₄ . How many times, on average, do you use USC Leavey collaborative workrooms per week?	0 – 10 or more
Q ₅ . How many meeting attendees, on average, do you have?	1 – 10 or more
Q ₆ . What is your average meeting time duration? (in hour)	1 – 5 or more
Q ₇ . How much do you consider energy savings while requesting scheduling meetings?	1 (Do not consider at all) – 7 (Extremely consider)
Q ₈ . I consider myself an environmentalist.	1 (Disagree) – 7 (Agree)

Table 9 Survey I: Questionnaire

Assumption (A)	Let us assume that you would like to schedule a meeting next week using the central meeting reservation system, which is currently used at USC Leavey library.
Question (Q)	Q ₁ . What is your preferred time range to start the meeting on each day of the week? (Note: Please consider your actual class and other meeting constraints while answering this.)
	Q ₂ . What locations do you prefer for your meeting among the rooms that you chosen? For your information, the number in the parentheses indicates the maximum capacity of each room. (Note: Please try to answer this based on your past experience at USC Leavey library.)

**Fig. 16** Diversity of people's flexibility

We analyzed their profile including the details of their meeting requests and their flexibility in terms of time and locations considering their real constraints. Tables 8 & 9 show a list of detailed questions in the questionnaire used during the survey.

Figure 16 shows the distribution of the time and location flexibility. The x-axis shows the discretized flexibility level and their corresponding frequency in percent-

age is provided on the y-axis. People reported varied levels of time and location flexibility. The average time flexibility (α^T) was 25.34% and the measured minimum and maximum time flexibility were 9.86% and 42.86%, respectively. The average location flexibility (α^L) was 16.05% and its range was 0 to 38.24%. This survey result clearly shows that people have fairly diverse flexibility levels and gives us the insight that there is a significant potential to conserve energy by exploiting scheduling flexibility in TESLA.

6.2 Survey for requested flexibility

We conducted a second survey to understand what types of feedback are most effective to change flexibility while scheduling meetings. We consider two test conditions: (i) feedback without motivation (Test Group I) (e.g., if necessary, do you think you will be able to provide more options in terms of time and location?), and (ii) feedback with motivation including average flexibility provided, and environmental motives (Test Group II) (e.g., on average, people who are using this system give 3–4 hour range for their available time on each day and 5–6 rooms for their available locations. This helps the system to compute more energy-efficient schedules that lead to energy savings by about 30% at the testbed building, which is equivalent to \$5,765 per year. Do you think you will be able to provide more options in terms of time and location?). A more detailed list of questions is shown in Table 10.

Hypothesis 1 *More informed feedback (provided to subjects in Test Group II) will be more effective to conserve energy than feedback without motivation (Test Group I).*

To test the hypothesis above, we recruited 22 students with the same requirement of the earlier survey. Subjects were randomly tested under two different conditions when they accessed the online survey, and each test group had 11 individuals respectively.

Table 11 shows the average flexibility change in percentage (0–100%) of two test groups. Thus, higher values indicate that more participants comply and increase their scheduling flexibility to higher levels. When we provided more informed feedback including environmental motives (Group II), participants tripled their flexibility increase percentage (17.12%). In Group I, participants only increase their flexibility level by 5.15% on average. The difference is statistically significant and provides strong evidence for the hypothesis (t-test; $p < 0.01$). This study shows that we can conserve energy by investigating methods to improve motivation to conserve energy by adjusting their flexibility.

In this trial study, we have learned that although occupants in commercial buildings do not have a direct financial incentive in saving energy, proper motivations can achieve a higher compliance rate for the energy-related suggestion with a specific focus on their flexibility. This study specifically gives us the insights that there is a significant potential to conserve energy by investigating effective and tailored methods to improve occupants' motivation to conserve energy while handling energy-efficient

Table 10 Survey II: Questionnaire

A.	Let us assume that you would like to schedule a meeting next week using the central meeting reservation system, which is currently used at USC Leavey library.
	Group I (Simple)
	Q ₁ . What is your preferred time range to start the meeting on each day of the week? (Note: Please consider your actual class and other meeting constraints while answering this.)
	Q ₂ . What locations do you prefer for your meeting among the rooms that you chosen? For your information, the number in the parentheses indicates the maximum capacity of each room. (Note: Please try to answer this based on your past experience at USC Leavey library.)
	On the previous page, you were asked about your preferred time and locations to schedule meetings. Given your choice, please answer following questions.
	Q ₃ . If necessary, do you think you will be able to provide more options in terms of time? If so, for each day of the week, what will be your extended available time range for your meeting? If you do not think you will be able to provide additional options, please skip this question. (Note: Please consider your actual class and other meeting constraints while answering this.)
	Q ₄ . Likewise, what additional locations would you consider for your meeting? If you do not think you will be able to provide additional options, please skip this question. For your information, the number in the parentheses indicates the maximum capacity of each room. (Note: Please try to answer this based on your past experience at USC Leavey library.)
	Group II (Complex)
	Q ₁ . What is your preferred time range to start the meeting on each day of the week? (Note: Please consider your actual class and other meeting constraints while answering this.)
	Q ₂ . What locations do you prefer for your meeting among the rooms that you chosen? For your information, the number in the parentheses indicates the maximum capacity of each room. (Note: Please try to answer this based on your past experience at USC Leavey library.)
Q.	On the previous page, you were asked about your preferred time and locations to schedule meetings. Given your choice, please answer following questions.
	Q ₃ . If necessary, do you think you will be able to provide more options in terms of time? If so, for each day of the week, what will be your extended available time range for your meeting? If you do not think you will be able to provide additional options, please skip this question.
	On average, people who are using this system give 3–4 hr range for their available time on each day. This helps the system to compute more energy-efficient schedules that lead to energy savings by about 30% at USC Leavey library, which is equivalent to \$5,765 per year.
	(Note: Please consider your actual class and other meeting constraints while answering this.)
	Q ₄ . Likewise, what additional locations would you consider for your meeting? If you do not think you will be able to provide additional options, please skip this question. For your information, the number in the parentheses indicates the maximum capacity of each room.
	On average, people who are using this system choose 5–6 rooms for their available locations. This helps the system to compute more energy-efficient schedules that lead to energy savings by about 20% at USC Leavey library, which is equivalent to \$3,845 per year.
	(Note: Please try to answer this based on your past experience at USC Leavey library.)

scheduling problems. However, at the same time, in order to deploy our TESLA system in the real-world while keeping people in the loop, we have a number of research challenges that have to be addressed. Most notably, in a commercial setup where people do not have a direct financial incentive to save energy, a different incentive mechanism to effectively motivate them and keep them as active participants in energy

Table 11 Flexibility manipulation with various feedback (%)

	Group I	Group II
Average amount of flexibility change	5.15	17.12

saving activities might potentially be required; determining the importance of such mechanisms or if they are needed in the first place is a topic for future work [1, 3, 8, 11, 45]. Over time, people will be able to observe the impact of their input (e.g., flexibility) while scheduling meetings and whether or not people engaged with TESLA on a day-to-day basis will provide flexibility to the extent they could remain to be determined. Thus, while this paper has provided a critical first step in flexibility-based energy savings, and provided algorithms to accomplish such savings, a future implementation will need to take the next step to investigate topics such as motivation and incentives.

7 Related Work

TESLA differs from previous work given its focus on: (i) commercial buildings (whereas significant previous work [4, 17, 27, 33, 36, 37, 46] has focused on residential buildings; (ii) comfort-based energy-efficient incremental scheduling using an SMILP; and (iii) identifying key meetings for effectively adjusting people’s scheduling flexibility. Furthermore, TESLA is evaluated on real meeting data (over 110,000 meetings and events) that have been collected from more than 500 rooms in nine educational buildings at USC and SMU. This combination of research contributions sets our work apart from previous research.

Energy Systems and Scheduling: Multiagent systems have been considered to provide sustainable energy for buildings and smart grid management. Stein *et al.* [37] introduced a novel online mechanism that schedules the allocation of an expiring and continuously-produced resource to self-interested agents with private preferences while focusing on the fairness using pre-commitment in the smart grid domain, which is not directly applicable in commercial buildings. Miller *et al.* [27] investigated how the optimal dispatch problem in the smart grid can be framed as a decentralized agent-based coordination problem and presented a novel decentralized message passing algorithm. Their work was empirically evaluated in large networks using real distribution network data. In addition, [17] addressed research challenges to optimally schedule charging plug-in Electric Vehicles (EVs) in the smart grid.

To model and optimize building energy consumption, Mamidi *et al.* [24] developed smart sensing and adaptive energy management agents to decrease energy consumptions by HVACs in buildings. They showed that in the educational building, these sensor agents can be used to accurately estimate the number of occupants in each room and predict future occupancy relying on machine learning to intelligently control HVAC systems. This work can be used for enhancing TESLA by incorporating user occupancy data into our system to more effectively determine key meetings. Ramchurn *et al.* [33] considered more complex deferrable loads and managing comfort in the residential buildings. There has been other work focusing on scheduling of

home appliances considering user preferences [4,36,46]. In particular, they consider inferred user's preferred usage profile while scheduling home appliances in residential buildings, which is considered as a fixed constraint. Our work is different as it does not only maximize energy savings while considering users' preferences, but also effectively interacts with users to change their flexibility to achieve further energy savings. More recently, there has been some work focusing on energy-aware scheduling in commercial buildings [23]. The authors only consider the HVAC systems and ignore other significant energy consumers such as lighting and electronics in commercial buildings while optimizing schedules based on the given *fixed* constraints. TESLA is different by focusing on an energy-oriented scheduling while considering major energy consumers (HVACs, lighting and electronics) together in commercial buildings. TESLA also identifies key meetings for flexibility change, an aspect that is missing in this previous work.

Wainer *et al.* [43] presented a set of protocols for scheduling a meeting among agents that represent their respective user's interests and evaluated the suggested protocols while handling meeting scheduling problems. The objective in their work is to find the optimal protocol to reach agreement among agents, which does not explicitly account for energy. In our own previous work [20,21], we consider meeting (re)location problems by exchanging messaging among agents. Although that work focused on minimizing energy consumption, it relied on the reactive scheduling and no flexibility model nor key meetings were considered.

In a multiagent community, there has been a significant amount of work that has focused on meeting/event scheduling based on the distributed constraint optimization (DCOP) formulation [22,40]. They provide distributed scheduling frameworks that are limited to dynamic scheduling problems. In addition, they focused on scheduling meetings without energy considerations. TESLA differs from their work as it explicitly aims to conserve energy while scheduling incrementally/dynamically arriving requests.

Online scheduling techniques have been investigated to handle incremental requests considering temporal flexibility [12,31]. Our work is different by focusing on energy-oriented scheduling in commercial buildings while allowing people to play a part in optimizing the operation in the building.

Social Influence in Human Subject Studies: We leverage insights from social psychology in understanding and designing reliable and accurate human behavior models. Wood and Neal [45] have studied the potential of interventions to reduce energy consumption and they have shown that it not only helps to change workplace energy consumption but also to establish energy use habits that maintain over time. Anderson *et al.* [3] have investigated social influence in energy use behavior, which can be used for enhancing TESLA to effectively change people's preferences and energy behaviors to conserve energy. Abrahmase *et al.* [1] reviewed 38 interventions aimed at reducing household energy consumption, and they concluded that normative feedback about energy use is the most promising strategy for reducing and maintaining low consumption. However, it focused on residential environments, which is different from our work. In a recent study, Carrico and Riemer [8] provided monthly normative feedback via email to occupants of a commercial building about their own buildings' energy use in comparison with other similar buildings. Faruqui *et al.* [11] reviewed

past experiments and pilot projects to evaluate the effect of in-home displays (IHDs) on energy consumption. Our work is different because we simultaneously consider multiple criteria including energy consumption *and* occupant comfort level.

In social psychology, there has been a significant amount of work to figure out the correlation between irritation/distraction factors and persuasion. McCullough and Ostrom [25] and Cacioppo and Petty [7] discussed that message repetition would increase positive attitudes in a situation where highly similar communications are used and showed that there is a positive relationship between the number of presentations and attitude from general social psychology perspectives. Focusing on a commercial advertisement, Pechmann and Stewart [30] predicted the effectiveness of different strategies on advertising and examined the effects of message repetition on attitude changes. In addition, Baron *et al.* [5] discussed that distractions affect behavior decisions, but they are more or less effective in increasing persuasion depending upon whether people can easily ignore the distraction. TESLA could benefit from these studies as they give us a deep understanding regarding an effective interaction mechanism design between TESLA and users.

8 Conclusion

This paper focused on energy savings in commercial buildings, and started with the observation that meetings play a significant role in this energy consumption. The key contribution of this paper is not just our agent TESLA, but also importantly, TESLA's analysis of real-world data — 32,000 meetings from the University of Southern California (USC), and 80,000 meetings from the Singapore Management University (SMU) — to show the power of flexibility. TESLA's promise of energy savings is rooted in this real data, and illustrates that significant energy savings comes not from imposing any complex interaction protocol on humans, but from the simple action of providing schedule flexibility. More specifically, this paper provided four key contributions. First, TESLA provided online scheduling algorithms to solve a stochastic mixed integer linear program (SMILP) while considering the diversity of people's flexibility for energy-efficient scheduling of incrementally/dynamically arriving meetings and events. Second, TESLA also included an algorithm to effectively identify key meetings that lead to significant energy savings by adjusting their flexibility. Third, this paper provides an explanation of why TESLA works by presenting extensive analysis on the energy savings achieved by TESLA. Lastly, surveys of real users were provided indicating that TESLA's savings can be realized in practice. We showed that, compared to the current systems, TESLA can substantially reduce the overall energy consumption. Although we have focused on evaluating TESLA in commercial buildings, TESLA can be applied to general scheduling domains where schedule flexibility plays a key role for conserving energy such as home appliance scheduling in residential buildings or resource scheduling for manufacturing systems.

Acknowledgements This material is based upon work supported by the National Science Foundation under Grant No. 1231001. Any opinions, findings, and conclusions or recommendations expressed in

this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

1. Abrahamse, W., Steg, L., Vlek, C., Rothengatter, T.: A review of intervention studies aimed at household energy conservation. *J Environ. Psychol.* **25**, 273–291 (2005)
2. Ahmed, S., Shapiro, A., Shapiro, E.: The sample average approximation method for stochastic programs with integer recourse. *SIAM Journal of Optimization* **12**, 479–502 (2002)
3. Anderson, K., Lee, S., Menassa, C.: Effect of social network type on building occupant energy use. In: *Buildsys*, pp. 17–24. ACM (2012)
4. Bapat, T., Sengupta, N., Ghai, S.K., Arya, V., Shrinivasan, Y.B., Seetharam, D.: User-sensitive scheduling of home appliances. In: *SIGCOMM* (2011)
5. Baron, R., Baron, P., Miller, N.: The relation between distraction and persuasion. *Psychological Bulletin* **80**(4), 310 (1973)
6. Beale, E.: On minimizing a convex function subject to linear inequalities. *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 173–184 (1955)
7. Cacioppo, J., Petty, R.: Effects of message repetition on argument processing, recall, and persuasion. *Basic and Applied Social Psychology* **10**(1), 3–12 (1989)
8. Carrico, A., Riemer, M.: Motivating energy conservation in the workplace: An evaluation of the use of group-level feedback and peer education. *J Environ. Psychol.* **31** (2011)
9. Dantzig, G.B.: Linear programming under uncertainty. *Management Science* **1**(3–4), 197–206 (1955)
10. Dubois, D., Fargier, H., Fortemps, P.: Fuzzy scheduling: Modelling flexible constraints vs. coping with incomplete knowledge. *European Journal of Operational Research* **147**(2), 231–252 (2003)
11. Faruqui, A., Sergici, S., Sharif, A.: The impact of informational feedback on energy consumption - a survey of the experimental evidence. *Energy* **35** (2010)
12. Gallagher, A., Zimmerman, T., Smith, S.: Incremental scheduling to maximize quality in a dynamic environment. In: *ICAPS* (2006)
13. Ghavamzadeh, M., Mahadevan, S., Makar, R.: Hierarchical multi-agent reinforcement learning. *Autonomous Agents and Multi-Agent Systems* **13**(2), 197–229 (2006)
14. Guestrin, C., Venkataraman, S., Koller, D.: Context-specific multiagent coordination and planning with factored mdps. In: *AAAI/IAAI*, pp. 253–259 (2002)
15. INFOCOM: Meetings in america: A study of trends, costs, and attitudes toward business travel and teleconferencing, and their impact on productivity. Whitepaper, INFOCOM (2001)
16. Kall, P., Wallace, S.W.: *Stochastic programming*. John Wiley and Sons Ltd (1994)
17. Kamboj, S., Kempton, W., Decker, K.S.: Deploying power grid-integrated electric vehicles as a multi-agent system. In: *AAMAS* (2011)
18. Kelso, J.D. (ed.): *Buildings Energy Data Book*. U.S. Dept. of Energy (2011)
19. Kwak, J., Varakantham, P., Maheswaran, R., Chang, Y.H., Tambe, M., Becerik-Gerber, B., Wood, W.: TESLA: An energy-saving agent that leverages schedule flexibility. In: *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2013)
20. Kwak, J., Varakantham, P., Maheswaran, R., Tambe, M., Hayes, T., Wood, W., Becerik-Gerber, B.: Towards robust multi-objective optimization under model uncertainty for energy conservation. In: *AAMAS Workshop on Agent Technologies for Energy Systems (ATES)* (2012)
21. Kwak, J., Varakantham, P., Maheswaran, R., Tambe, M., Jazizadeh, F., Kavulya, G., Klein, L., Becerik-Gerber, B., Hayes, T., Wood, W.: SAVES: A sustainable multiagent application to conserve building energy considering occupants. In: *AAMAS* (2012)
22. Maheswaran, R.T., Tambe, M., Bowring, E., Pearce, J.P., Varakantham, P.: Taking dcop to the real world: Efficient complete solutions for distributed multi-event scheduling. In: *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pp. 310–317. IEEE Computer Society (2004)
23. Majumdar, A., Albonese, D.H., Bose, P.: Energy-aware meeting scheduling algorithms for smart buildings. In: *Buildsys*, pp. 161–168. ACM (2012)
24. Mamidi, S., Chang, Y.H., Maheswaran, R.: Improving building energy efficiency with a network of sensing, learning and prediction agents. In: *AAMAS* (2012)
25. McCullough, J., Ostrom, T.: Repetition of highly similar messages and attitude change. *Journal of Applied Psychology* **59**(3), 395 (1974)

26. Michigan State University: New classroom scheduling methods save energy, money for msu. <http://news.msu.edu/story/6501/> (2009)
27. Miller, S., Ramchurn, S.D., Rogers, A.: Optimal decentralised dispatch of embedded generation in the smart grid. In: AAMAS (2012)
28. Mohsenian-Rad, A.H., Leon-Garcia, A.: Optimal residential load control with price prediction in real-time electricity pricing environments. *Smart Grid, IEEE Transaction on* **1**(2), 120–133 (2010)
29. Pagnoncelli, B., Ahmed, S., Shapiro, A.: Sample average approximation method for chance constrained programming: theory and applications. *Journal of optimization theory and applications* **142**(2), 399–416 (2009)
30. Pechmann, C., Stewart, D.: Advertising repetition: A critical review of wearin and wearout. *Current issues and research in advertising* (1988)
31. Policella, N., Smith, S.F., Cesta, A., Oddi, A.: Incremental scheduling to maximize quality in a dynamic environment. In: ICAPS (2004)
32. Portland State University: Efficient class scheduling conserves energy. <http://goo.gl/cZwgB> (2012)
33. Ramchurn, S.D., Vytelingum, P., Rogers, A., Jennings, N.R.: Agent-based control for decentralised demand side management in the smart grid. In: AAMAS (2011)
34. Scerri, P., Pynadath, D.V., Tambe, M.: Towards adjustable autonomy for the real world. *JAIR* **17**, 171–228 (2002)
35. Shapiro, A., Dentcheva, D., Ruszczyński, A.: Lectures on stochastic programming: modeling and theory, vol. 9. Society for Industrial Mathematics (2009)
36. Sou, K.C., Weimer, J., Sandberg, H., Johansson, K.H.: Scheduling smart home appliances using mixed integer linear programming. In: CDC-ECC (2011)
37. Stein, S., Gerding, E., Robu, V., Jennings, N.: A model-based online mechanism with pre-commitment and its application to electric vehicle charging. In: AAMAS (2012). URL <http://eprints.soton.ac.uk/273082/>
38. Strbac, G.: Demand side management: Benefits and challenges. *Energy Policy* **36**(12), 4419–4426 (2008)
39. Subramanyam, S., Askin, R.G.: An expert systems approach to scheduling in flexible manufacturing systems. Ph.D. thesis, University of Iowa (1985)
40. Sultanik, E., Modi, P.J., Regli, W.C.: On modeling multiagent task scheduling as a distributed constraint optimization problem. In: Proceedings of the 20th International Joint Conference on Artificial Intelligence, pp. 1531–1536 (2007)
41. U.S. Department of Labor: Average energy prices in the los angeles area. http://www.bls.gov/ro9/cpilosa_energy.htm (2012)
42. Varakantham, P., Kwak, J., Taylor, M.E., Marecki, J., Scerri, P., Tambe, M.: Exploiting coordination locales in distributed pomdps via social model shaping. In: ICAPS (2009)
43. Wainer, J., Jr., P.R.F., Constantino, E.R.: Scheduling meetings through multi-agent negotiations. *Decision Support Systems* **44**(1) (2007)
44. Wang, C., de Groot, M., Marendy, P.: A service-oriented system for optimizing residential energy use. In: Web Services, IEEE International Conference on (2009)
45. Wood, W., Neal, D.: The habitual consumer. *Journal of Consumer Psychology* **19**, 579–592 (2009)
46. Xiong, G., Chen, C., Kishore, S., Yener, A.: Smart (in-home) power scheduling for demand response on the smart grid. In: ISGT (2011)