Towards Optimal Planning for Distributed Coordination Under Uncertainty in Energy Domains

Jun-young Kwak, Pradeep Varakantham*, Milind Tambe, Laura Klein, Farrokh Jazizadeh, Geoffrey Kavulya, Burcin B. Gerber, David J. Gerber
University of Southern California, Los Angeles, CA, 90089
*Singapore Management University, Singapore, 178902
{junyounk,tambe,lauraakl,jazizade,kavulya,becerik,dgerber}@usc.edu,
*pradeepv@smu.edu.sg

ABSTRACT

Recent years have seen a rise of interest in the deployment of multiagent systems in energy domains that inherently have uncertain and dynamic environments with limited resources. In such domains, the key challenge is to minimize the energy consumption while satisfying the comfort level of occupants in the buildings under uncertainty (regarding agent negotiation actions). As human agents begin to interact with complex building systems as a collaborative team, it becomes crucial that the resulting multiagent teams reason about coordination under such uncertainty to optimize multiple metrics, which have not been systematically considered in previous literature. This paper presents a novel multiagent system based on distributed coordination reasoning under uncertainty for sustainability called SAVES. There are three key ideas in SAVES: (i) it explicitly considers uncertainty while reasoning about coordination in a distributed manner relying on MDPs; (ii) human behaviors and their occupancy preferences are incorporated into planning and modeled as part of the system; and (iii) the influence of various control strategies for multiagent teams is evaluated on an existing university building as the practical research testbed with actual energy consumption data. We empirically show the preliminary results that our intelligent control strategies substantially reduce the overall energy consumption in the actual simulation testbed compared to the existing control means while achieving comparable average satisfaction level of occupants.

Categories and Subject Descriptors

I.2.11 [**ARTIFICIAL INTELLIGENCE**]: Distributed Artificial Intelligence

General Terms

Algorithms, Human Factors

Keywords

Sustainability, Multi-Objective Optimization, Energy, Satisfaction, Multiagent Systems

1. INTRODUCTION

Over the decades, energy issues have been getting more important. In the U.S., about 40% of energy consumption is from buildings (shown in Figure 1), of which 25% is associated with heating and cooling [21] at an annual cost of $40 billion [21]. Furthermore, on an annual basis, buildings in the United States consume 73% of its electricity. Recent developments in multiagent systems are opening up the possibility of deploying multiagent teams to achieve complex goals in such energy domains that inherently have uncertain and dynamic environments with limited resources.

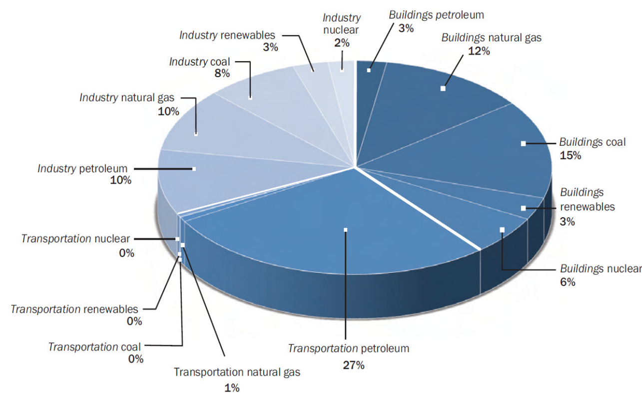

Figure 1: Distribution of US energy use in 2006, grouped by end-use sector (transportation, buildings, industry). Annual consumption for 2007 was 101.6 quads (10^{15} BTU)

This paper focuses on a novel planning method for distributed coordination under uncertainty (regarding agent negotiation actions) to optimize multiple competitive objectives: i) amount of energy used in the buildings; ii) occupant's comfort level; and iii) practical usage considerations. There have been some trials to balance energy consumption and enhancement of building services and comfort levels [15, 19, 23, 25] and to monitor and collect energy consumption data [15, 16] in energy domains. Other works have explicitly focused on design optimization and use of multiagent systems [13, 17, 20] in different domains. In addition, some multiagent systems [5, 6, 8, 9, 10, 11, 22] and the underlying theory for their decision supports [12] have been employed to model home automation systems. Unfortunately, past work in the energy domain has three key weaknesses. First, they do not consider uncertainty while reasoning about coordination and mostly rely on deterministic plans. Second, they limitedly incorporate intelligence of occupancy or occupancy preferences into the system and thus occupants are not explicitly modeled as agents in the system. Third, their works are mostly evaluated in their own simulation environments, which are not constructed on the actual energy data and occupants' responses in the buildings. Thus, their assumptions may not be realized in real-world problems.

This paper presents a novel multiagent system based on distributed coordination reasoning under uncertainty for sustainability called SAVES (Sustainable multi-Agent systems for optimiz-

Figure 2: Testbed - Educational Building at USC

ing Variable objectives including Energy and Satisfaction). SAVES provides three key contributions to overcome limitations in past work. First, we explicitly consider uncertainty while reasoning about coordination in a distributed manner. In particular, we rely on MDPs (Markov Decision Problems) to model agent interactions, specifically focusing on rescheduling meetings, which will be extended to decentralized MDPs. Second, human behaviors and their occupancy preferences are incorporated into planning and modeled as part of the system. As a result, SAVES is capable of generating an optimal plan not only for building usage but also for occupants. Third, the influence of various control strategies for multiagent teams is evaluated on an existing university building as the practical research testbed with actual energy consumption data in the simulation. Since the simulation environment is based on actual data, this result can be easily deployed into the real-world. Preliminary results show that our intelligent control strategies substantially reduce the overall energy consumption in the actual simulation testbed compared to the existing control means while achieving comparable average satisfaction level of occupants.

2. MOTIVATING DOMAINS

This work is motivated by energy domains where multiagent coordination can be the key issue. To pin down the domain problem, we consider an actual educational building (shown in Figure 2) as a representative test case to measure and collect the energy consumption and responses of occupants because it is a multi-functional building of sufficient size and activity for research.[1] Furthermore, the building is representative in that it has been designed with a building management system, and it provides a good environment to test various control strategies to mitigate energy consumption. The research can be easily generalized to other building types, where we can observe many different types of energy-use awareness based on the behavioral patterns of occupants in the buildings.

Our research testbed is focused on testing different operation optimization strategies based on the scope of occupant behaviors and schedules. The simulation component will include the building, its human occupants, and its facility management. It will then interact with the occupants and management via proxy agents [24] to advise them on how to reduce energy use while measuring occupant comfort level. More specifically, human occupants are divided into two main categories — permanent and temporary building occupants. Permanent building occupants include office resident such as faculty, staff, and researchers and laboratory residents like researchers in web labs, structural labs, etc. Temporary building occupants have scheduled occupants who include students or faculties attending classes or meetings and unscheduled occupants who are students or faculty using common lounge or dining space. In this domain, proposed human energy behaviors are entering/leaving a room, turning on/off light sources, turning on/off computers and other electronics, adjusting thermostat (heating or cooling), adjusting window shading, opening an operable door or window, adjusting personal clothing, and adjusting activity level. Building components and equipments that are another type of agent in the buildings include HVAC systems, which are composed of air handler units, VAV boxes, temperature sensors, and thermostats, lighting systems, office electronic devices such as computer and AV equipments, and laboratory equipments. Measurement of energy consumption for each equipment action may be estimated from design specifications. In our work, we choose and implement a subset of agents and their energy-related behaviors listed above.

[1]The size and other parameters of the building are given in the evaluation section.

3. RELATED WORK

With rising energy costs, the need to design and integrate scalable energy consumption reduction strategies in buildings calls for novel approaches. There are numerous challenges associated with energy resources such as supply and depletion of energy resources and heavy environmental impacts [19] (ozone layer depletion, global warming, climate change, etc.). The rise in energy consumption in buildings can be attributed to several factors such as enhancement of building services and comfort levels [15, 19, 23, 25], through heating, cooling and lighting needs and increased time spent indoors [19].

To model and optimize buildings' energy consumption, building owners and facility managers are demanding robust, intelligent and adaptable performance monitoring techniques. These techniques are important in energy consumption data collection [15, 16] and ambient environmental conditions control [15]. Existing heating, cooling, ventilation, and lighting systems generally operate with no intelligence of occupancy or occupancy preferences and therefore are unable to optimize operations. Even more, no feedback is available to occupants about how their actions and schedules impact building energy consumption. To realize both tangible benefits such as energy and operation savings, value property, reduction in occupant complaints as well as the intangible benefits such as occupant comfort and satisfaction, designers must develop energy adaptive capabilities within the building environmental control systems.

Abras *et al.* [5], Conte *et al.* [9] and Roy *et al.* [22] have employed multiagent systems to model home automation systems (or smart homes) and simulating control algorithms to evaluate performance. While there is relevance in terms of the problem domain and employing multiagent systems, our representation and approaches are different in having to account for human preferences and decisions directly.

Research by Fahrioglu *et al.* [14], Mohsenian-Rad *et al.* [18] and Caron *et al.* [7] provide incentive compatible mechanisms for distribution of energy among interested parties. This thread of research is complementary, especially in designing incentives for humans to reveal their true energy preferences. However, these approaches assume a centralized controller with whom all the members interact, which is not present in our domain. Instead, there are peer-to-peer negotiations between humans regarding their energy consumption and comfort level.

Figure 3: Overall System Design

4. DESIGN DECISIONS

The SAVES system consists of a simulation module, an input/output module to communicate with agents, and an underlying reasoning and planning module. Figure 3 shows a generic loop of the system. In particular, the input/output module first collects data and constructs the world model. Given the world model, the reasoning and planning module generates policies to achieve the given objectives in the context of coordination. With these world model and generated policies, the simulation module models agents' physical and behavioral interactions in the system and realize the coordination in the actual world via the input/output module. We now describe the modules as well as the particular instantiations of these modules in the energy domain.

The simulation module provides a 2D, OpenGL environment based on the open-source project OpenSteer [2] as shown in Figure 4(a) & 4(b). The simulation module consists of two different types of agents as described below, modeling their physical and behavioral interactions. It can be used for efficient statistical analysis of different control strategies in buildings before deploying the system to an actual physical world.

The input/output module makes a connection among different modules in the system by collecting actual data in the domain, transferring data to the reasoning module, sending output results to either the simulation or deployed module in the world to represent outputs, and providing means to communicate with agents via proxy and handheld devices.

The coordination and planning module generates optimal policies to achieve the given team missions considering multiagent interactions in the context of coordination in the mutiagent setup.

Here we describe the design issues regarding agents, first introducing building component agents and human agents, then detailing the method to calculate the properties of agents, and finally discussing different control strategies considering agent interactions.

4.1 Building Component Agents

We consider three building component agent categories: a HVAC (Heating, Ventilating, and Air Conditioning) agent, a lighting agent, and an appliance agent. The HVAC agent (Figure 5(a)) is modeled based on the principles of thermodynamics, fluid mechanics, and heat transfer. We assume that this agent mainly controls the temperature of the assigned zone. The lighting agent (Figure 5(b))

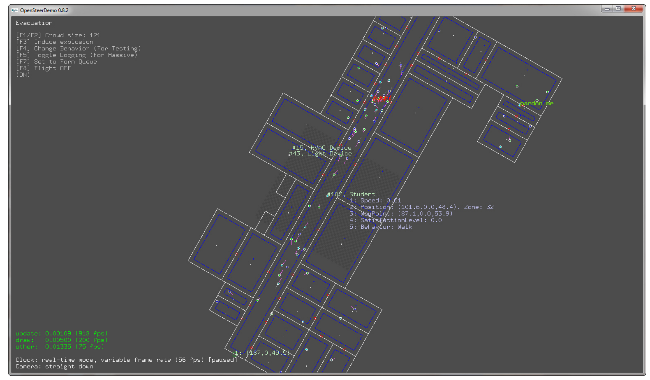

(a)

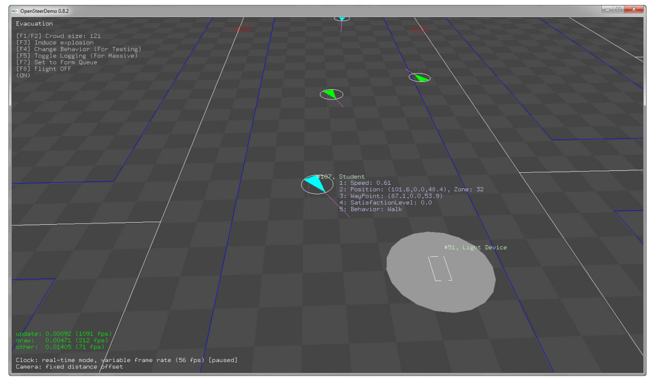

(b)

Figure 4: Screen capture

controls the lighting level of the room. For the appliance agent, we only include the computer device including the desktop and laptop computers in this work. These agents have two possible actions: "on" [2] and "standby". When the lighting or appliance agents are "on", they consume some fixed amount of energy. We measure the average amount of energy used by these agents, which will be detailed in Section 6.

Since the energy consumption of HVAC agents relies on a set of parameters including the temperature change in the space, air flow, and number of people, etc., the average value cannot be simply measured. Instead, we describe how to compute the energy use by HVAC agents below.

[2]Note that the "on" action can be divided into several actions with different output levels, e.g., "on" with 30%, 60%, and 100% power.

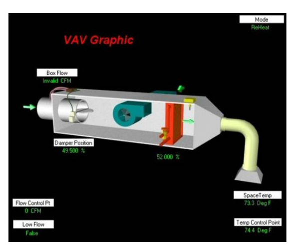

(a) HVAC Agent

(b) Lighting Agent

Figure 5: Agents

Calculating Total Energy Consumption: Since the building is composed of a large number of HVACs and they are the main consumers of the energy, it is important to choose the right set of parameters and reasonable values for them. In particular, the energy consumption of HVAC agents is calculated as following [1, 4] mainly based on changes in air temperature and airflow speeds:

$$Q = \frac{1.1 \times CFM \times \Delta T}{3412.3},$$

$$\Delta T = log(\frac{CFM}{C}),$$

where Q is the amount of energy used (kWh), CFM is an air volume flow rate (ft^3/min), which is typically ranged between 500–1500 (ft^3/min), ΔT is the temperature change in a zone (°F), and C is a scale factor.

4.2 Human Agents

There are four different types of human agents such as a faculty, staff, graduate student, and undergraduate student. Each agent has access to a subset of the six available behaviors according to their types — wander, attend the class, go to the meeting, teach, study, and perform research, any one of which may be active at a given time, where the behavior is selected via the given class and meeting schedules.

During execution of these behaviors, individual travelers may move at integer speeds from 0 to 3. Each agent also has specific levels of emotions and information about the environment. Specifically, every agent has a property about the satisfaction level based on the current environmental condition and knows his or her current location without any noise. A more extended discussion of the satisfaction property will take place below.

Calculating Satisfaction Level: The satisfaction level (SL) of an individual human agent is modeled as a percentage value between 0 and 100 (0 is fully unsatisfied, 100 is fully satisfied). SL of the individual occupant is calculated as following:

$$SL = 100.0 - PPD,$$

$$PPD = 100.0 - 95.0 \cdot \exp^{-(0.03353 \cdot PMV^4 + 0.2179 \cdot PMV^2)},$$

where SL is the satisfaction level (%), PPD is the Predicted Percent Dissatisfied (%), PMV is the Predicted Mean Vote. The PMV index is calculated from an equation of thermal balance for the human body in ASHRAE Standard [3], involving the parameter values shown in Table 1.

The PMV model uses heat balance principles to relate the seven key factors for thermal comfort listed in Table 1 to the average response of people on the above scale. The PPD index is calculated using the PMV as defined in [3]. It is based on the assumption that people respond about their comfort level with a number between -3 and +3 (-3 is cold, +3 is hot and 0 is neutral) on the thermal sensation scale and that the simplification that PPD is symmetric around a neutral PMV.

5. CONTROL STRATEGIES

In a given scenario, all agents within the simulation will use the same strategy. Possible strategies include: i) manual control strategy, ii) reactive control strategy, iii) proactive control strategy, and iv) proactive control strategy based on multiagent coordination.

5.1 Manual Control

Table 1: Parameters for the Satisfaction Level

Parameter	Value Range
Clothing	0.5 – 1.0 (light to heavy clothing)
Metabolic Rate	1.0 – 2.0 (low to high activity)
External Work	0
Air Temperature	20 – 28 (°C)
Radiant Temperature	20 – 28 (°C)
Air Velocity	0 – 0.2 m/s
Relative Humidity	30 – 60 %

The manual control strategy simulates the current building control strategy maintained by USC facility managers. Specifically, we assume that HVAC agents are not controlled by human agents and that appropriate temperature points are centrally set/given by facility managers. For HVAC agents, the CFM values are fixed throughout the simulation. In this control setting, HVAC agents always try to reach the pre-set temperature using the fixed CFM value regardless of the presence of human agents in the specific space and their preferences in terms of temperature. Lighting agents are controlled by only human agents. Control actions (i.e., turning on/off the light) of human agents are either deterministic or stochastic according to the type of action. In particular, when human agents enter the space, they always turn on the light. When they leave the space, they stochastically turn off the light. For appliance agents, we simply assume that they are always on.

5.2 Reactive Control

Since the manual control strategy simply follows the pre-defined policy provided by the facility managers, it is fairly easy to come up with action plans of building component agents. However, it does not adapt the given policies based on actual schedules or preferences of occupants in the building, and thus the building component agents are limited to adapt their control policies appropriately according to the dynamic changes. Particularly, HVAC agents keep operating to reach the desired point, even though the space is empty, which ends-up wasting energy. At the same time, since they do not consider occupants' preferences in the space and instead prioritize the pre-determined points, the average satisfaction level of occupants can decrease.

Here we discuss about another control strategy that building component agents reactively respond to the behaviors of human agents. In this setting, we assume that HVAC agents are not controlled by human agents and that appropriate temperature points are measured based on the average preference of human agents in the specific space. HVAC agents automatically turn on and off according to the presence of people and temperature set points, and the CFM values are adjusted based on the desired temperature point. In the reactive control strategy, the lighting and appliance agents are now automatically controlled. In particular, they are turned on and off according to the presence of people. For instance, when people enter the specific room, the lighting and appliance agents are automatically turned on, and when people leave the room, they are turned off.

While human agents follow their given schedules, with the reactive setting, the building component agents can act more intelligently than the manual policy as they operate based on human agents' actual needs. As a result, we can reduce the cases where the energy is wasted for unnecessary spaces, which will contribute to the reduction of the overall energy consumption.

5.3 Proactive Control

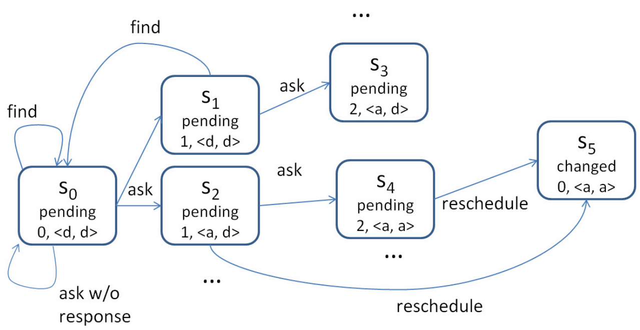

Figure 6: Simplified MDP model — d: disagree, a: agree

Although the reactive control strategy can adapt their policies based on actual needs of occupants in the building, this approach is still limited in a sense of optimality. In practice, there is a delayed effect in changing temperature. In other words, HVAC agents can only change a certain amount of degree in temperature per hour. This property exposes the weakness of the reactive control strategy. Although HVAC agents know the desired temperature of human agents at a specific time point, it takes a certain amount of time to reach the desired temperature point from the current air temperature, and the satisfaction level of occupants in the space will decrease during that time.

To overcome limitations of the reactive setting, we suggest a third control strategy operated in a proactive manner. Given the meeting and class schedules of human agents, the building component agents can predict: i) what their preferences are in terms of temperature, ii) how long it will take to reach the preferred temperature point from the current air temperature, iii) what CFM value is required, etc. In this setting, the building component agents can access the meeting/class schedules of human agents. Based on that prior knowledge, they now generate more optimal policies to reduce the overall energy consumption while maximizing the average satisfaction level of occupants in the building. For instance, a HVAC agent knows that the current air temperature is 55°F and the preferred temperature of the group of human agents who will use the space in 2.5 hours is 70°F. If the maximum possible temperature change by HVAC is 10°F/hr with the maximum CFM value, the HVAC agent predicts that it needs to change the temperate by 6°F per hour with a smaller CFM value which will use less energy. With this proactive plan, when human agents get to the space, the air temperature is already 70°F, which is the desired temperature point of people, and thus their satisfaction level increases.

5.4 Modeling Multiagent Coordination: MDP representation

With the existence of human agents, agent interactions are a fundamental aspect of our energy simulation. In SAVES , all agents share a common architecture based on MDP (Markov Decision Problem) frameworks, possessing varying degrees of knowledge about the world and other building agents (i.e., local knowledge).

MDPs have been used to tackle such real-world multiagent planning and coordination problems under transition uncertainty, which are described by a tuple $\langle S, A, T, R\rangle$, where

- $S = \{s_1, ..., s_k\}$ is a finite set of world states.
- A is the finite set of actions of an agent.
- $T : S \times A \times S \mapsto \mathbb{R}$ is the transition probability function, where $T(s'|s, a)$ denotes the transition probability from s to s' if an action a is executed.
- $R : S\times A\times S \mapsto \mathbb{R}$ is the reward function, where $R(s, a, s')$ denotes the reward that an agent gets by taking a from s and reaching s'.

We denote a policy computed by MDP $\pi : S \mapsto A$ is a mapping from world state to action. Our goal is effective multiagent team coordinations to minimize the total energy consumption while maximizing occupant's comfort level.

This section describes our MDP representation in the energy domain for illustration. The MDP model represents a class of MDPs covering all types of meetings for which the agent may take rescheduling actions. In our work, we construct a MDP for each meeting as shown in Figure 6. Alternatively, we can model all meetings in the building as a single MDP. However, if we consider a gigantic MDP model for rescheduling all meetings together, the number of states and actions exponentially explodes as the number of agents increases. In addition, the complexity to handle all possible coordinations among agents significantly increases, which is burdensome to handle within a reasonable amount of time.

As preliminary work, we construct a simplified MDP model for rescheduling meetings. For each meeting, a meeting agent can perform any of three actions — reschedule, find another slot [3], and ask. For the "ask" action [4], an agent can autonomously reduce its own autonomy and ask a human agent whether he or she agrees with rescheduling the existing meeting. The human agent can respond to the meeting agent with "agree" or "disagree".

The agent may choose to perform any of these actions in various states of the world. State is composed of three features: $\langle f_1, f_2, f_3\rangle$, where f_1 is the status whether meeting location and time is changed (i.e., pending or changed), f_2 is the number of "ask" actions invoked so far, and f_3 is a set of responses from all meeting attendees: $\langle rp_{i,1}, rp_{i,2}, ..., rp_{i,n}\rangle$, where n is the number of attendees of meeting i and $rp_{i,k}$ is a response of agent k to rescheduling meeting i (i.e., agree or disagree).

The MDP's reward function has its maximum value when the meeting agent invokes the "reschedule" action in the state where all meeting attendees agreed to reschedule. We denote the component of the reward function that focuses on the expected energy gain by rescheduling the meeting as r_{energy}. However, there is clearly a high team cost incurred by forcing all of the attendees to rearrange their schedules. This team cost is incorporated into the MDP's reward function by adding a negative reward, $r_{rearrange}$. The magnitude is also an increasing function in the number of attendees (e.g., rescheduling a meeting of a large group is more costly than rescheduling a one-on-one meeting). The overall reward function for taking the "reschedule" action, $a_{reschedule}$, in a state s is a weighted sum of these components:

$$R(a_{reschedule}, s) = \alpha \cdot r_{energy} + (1 - \alpha) \cdot r_{rearrange}$$

, where $0 \leq \alpha \leq 1$. In addition, a small amount of cost is incurred to invoke actions of "ask" and "find another slot".

The MDP's transition probabilities represent the likelihood over possible action outcomes. Specifically, the transition function is defined considering four factors: i) meeting constraints of attendees; ii) level of energy consciousness, which determines how much they care about energy; iii) degree of intimacy among occupants; and iv)

[3]This action can be modeled differently, e.g., delay 1 hr, delay 2 hrs, ..., delay n hrs, delay 1 day, ..., delay n days, cancel the meeting, change location (but same time), etc.

[4]The ask action can be divided into several actions with different amount of incentives, e.g., ask with 5% incentive, ask with 10% incentive, etc.

Table 2: Parameter Values for the Experiments

	Temperature	CFM	Likelihood
Manual	65–70°F	1500.0	50%
Reactive	Preference	500.0–1500.0	Automatic
Proactive	Preference	500.0–1500.0	Automatic
Proactive w/ MDP	Preference	500.0–1500.0	Automatic

Table 3: Parameter Values for the SL Calculation

Parameter	Value
Clothing	1.0
Metabolic Rate	1.2
External Work	0
Air Temperature	Zone temperature
Radiant Temperature	65°F
Air Velocity	0.1 m/s
Relative Humidity	40 %

the current status of responses, which can be related to emotional contagion within the group. Since we store the current set of responses from individual agents and number of "ask" actions called so far, the repeated "ask" action may result in different transitions. In particular, the "ask" action, by which the agent queries the human agent, has $2^{n_i} + 1$ possible outcomes, where n is the number of attendees of the meeting i. First, the human agent may not respond at all, in which case, the agent is performing the equivalent of a "wait" action for a given timeout. Other set of possible outcomes are decided depending on responses of meeting attendees as illustrated in Figure 6. We assume that the "find" action reset values of features in the state.

One possible policy, generated for a subclass of possible meetings, specifies "ask" and then "wait" in state S_0 of Figure 6, which prompts the agent to give up its autonomy. If the agent then reaches state S_1, the policy specifies "find", so the meeting agent figures out another available location or time for rescheduling. However, if the agent then reaches state S_2, the policy again chooses "ask", which asks the human agents once more to collect their responses. Similarly, if the agent reaches S_4, the "reschedule" action is chosen according to the policy.

Based on this MDP model, the agent reasons about different tradeoffs in team costs. This reasoning follows a fundamental tenet of teamwork in our system, that the individual team members act responsibly towards the team.

6. EMPIRICAL VALIDATION

We evaluate the performance of SAVES in our energy domain

Table 4: Average Energy Consumption of Lighting & Appliance Agents

Agent Type	Category	On	Off/Standby
Lighting	Office	0.128 kW/hr	0 kW/hr
	Conference room	0.192 kW/hr	0 kW/hr
	Classroom	0.768 kW/hr	0 kW/hr
Appliance	Desktop	0.150 kW/hr	0.010 kW/hr
	Laptop	0.050 kW/hr	0.005 kW/hr

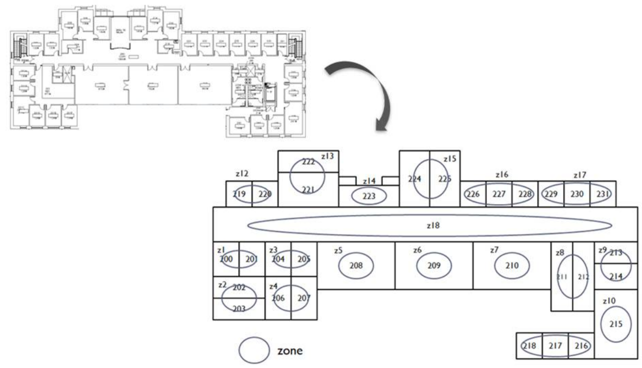

Figure 7: **Floor Plan - Educational Building at USC**

and compare four different control techniques: 1) manual control, 2) reactive control, 3) proactive control, and 4) proactive control with MDP. We focus on measuring two different criteria — total energy consumption (kWh) and average satisfaction level of occupants (%). The parameter values used in the experiments are shown in Table 2. In Table 2, column 2 shows the desired temperature for HVAC agents (Note: Preference in rows 3–5 means that the desired temperature is decided based on the average preference values of building occupants) and column 4 displays the likelihood value for the "turn off" action for the lighting agent. To calculate the energy consumption by the HVAC agent, we set the scale factor to 100.0. For the satisfaction calculation, we used the same parameter values in Table 3 while performing the experiments across four different control strategies. The experiments were run on Intel Core2 Quad-core 2.4GHz CPU with 3GB main memory. All techniques were evaluated for 100 independent trials throughout this section. We report the average values.

6.1 Experimental Domain Description

We have identified an educational building in conjunction with USC Facilities Management Services, as our practical testbed. This campus building is composed of classrooms, offices for faculty and staff, and conference rooms for meetings. Specifically, we use one floor of the actual university building in the experiments, which has 18 zones and 33 rooms as illustrated in Figure 7. There is one HVAC agent for each zone, and one lighting agent for each room. We also assume that each person in the office has either one desktop or laptop computer, and conference room and class room has two computers, respectively. There are four human agent categories: faculty, staff, graduate student and undergraduate student. Throughout the entire simulation, we consider a typical winter season in southern California (i.e., starting indoor temperature is 55°F in the simulation). During the simulation, indoor temperature goes down by -1°F per timestep, where each time step is 30 minutes. Possible temperature range in the building is between 50 and 90°F. Students follow 2010 Fall class schedule, and we generated the arbitrary meeting schedules for faculty, staff, and student agents. The measurement is performed during a working hour (i.e., 8:00am – 7:00pm), and the preference value of each occupant in temperature is randomly drawn from the uniform distribution between 60 – 70°F. To calculate the energy consumption of the lighting and appliance agents, we collect actual energy consumption data in the testbed building and used the average values shown in Table 4.

6.2 Comparison: Total Energy Consumption

We compared the cumulative total energy consumptions mea-

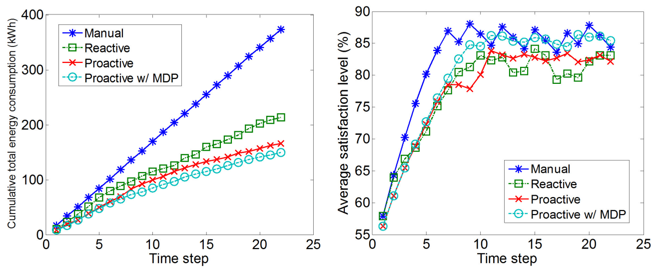

(a) Total Energy Consumption (b) Average Satisfaction Level

Figure 8: Comparison

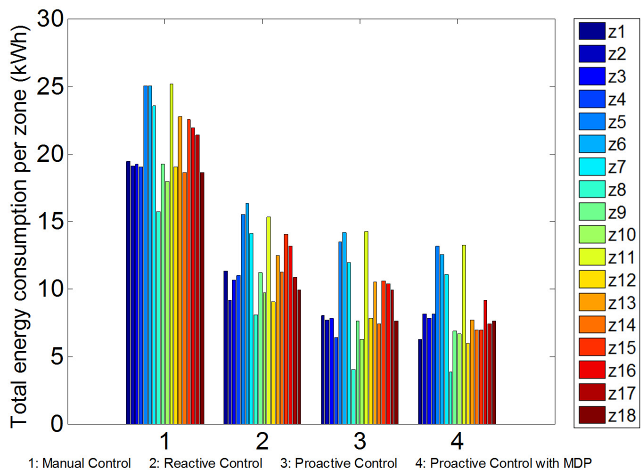

Figure 9: Energy Consumption Distribution

sured during work hours for all control strategies in the energy domain. Figure 8(a) shows the cumulative total energy consumption on the y-axis in *kWh* and the time step on the x-axis. Time step 1 indicates 8:00am and each time step increases by 30 minutes. As shown in the figure, the manual control strategy showed the worst result since it does not take into account behaviors or schedules of human agents and building component agents simply follow the predefined policies. The reactive and proactive control strategies showed lower energy consumptions than the manual setting by 43.0% and 55.6%, respectively. The proactive control strategy with the MDP model showed the best results among all different control strategies and statistically significant improvements (via t-tests) in terms of energy used in the testbed building, relative to other control strategies. Specifically, the proactive control with MDP reduced the energy consumption by 59.9% than the manual control strategy.

Although we did not tune the parameter values and only applied the simplified MDP model, with considering multiagent coordination in SAVES, we could achieve significant improvements. These outcomes are still preliminary results and yet only tested in the simulation environment, all experimental results were measured based on the actual data and testbed. Later, we will be able to show even more improvement with the optimally tuned parameters and extend our work to deploy it into the actual building with proxy agents. Furthermore, as we revise the equations shown in Section 4.1, we will be able to get more exact results for analysis.

Now, we analyze how various control strategies can cause different results. Figure 9 shows the energy consumption distribution over zones for all control strategies. In the figures, the x-axis shows the group number of data obtained by each control strategy and the y-axis displays the total energy consumption for each zone in *kWh*. The floor plan we used in the simulation has four different types of zones, which decides the total energy consumptions. Specifically, zones 1–4 (blue), 9 (green), and 12 (yellow) have two offices per zone, zones 5–7 (light blue or cyan) are class rooms, zones 13–15 (orange or red) are conference rooms, and zones 11 (yellow), 16 (light red), and 17 (red) have three offices per zone. As shown in the first group of Figure 9, the manual control strategy results in the similar level of energy consumptions according to the different types of zones. This result clearly indicates that the manual setting is only impacted by the physical constraints of the building space itself, which never considers the interactions among agents. The normalized standard deviation was 0.134. In the reactive (the second group in Figure 9) and proactive setting (the third group in Figure 9), it now started showing the difference in terms of the amount of energy used even within the same type of zones since those methods consider the actual behavioral patterns and schedules of human agents, and building component agents respond and adapt their policies based on them. Their normalized standard deviations are 0.205 and 0.312, respectively, which are higher than the value of the manual setting. Lastly, the proactive control strategy with the MDP model considers rescheduling of meetings. The target meetings to reschedule are ones with less than 4 people in the conference rooms in zones 13–15. We only considered the location reallocation and did not assume the meeting time can be also changed. New candidate locations are small faculty offices in zones 1–4. As shown in the fourth group of Figure 9, it showed increased energy consumptions in zones 1–4 due to the reallocated meetings, but simultaneously showed much more reduction in zones 13–15, as a result the overall energy consumption decreased. The normalized standard deviation was 0.313, which was the highest among different control strategies. These results give us a lesson that multiagent coordination/negotiation can benefit our model in SAVES, and by considering higher degree of coordination among agents, we will be able to achieve the significant energy reduction in this domain.

6.3 Comparison: Average Satisfaction Level

Here, we compare the average satisfaction level of human agents under different control strategies in the simulation. We used the equations discussed in Section 4.2.

Figure 8(b) shows the average satisfaction level in percentage on the y-axis and time step on the x-axis, which are the same as mentioned in the previous section.[5] As shown in the results, all methods were able to achieve at least 80% or higher results on average, and the manual and proactive with MDP settings showed the best results among them. Note that the equations to calculate the individual satisfaction level are based on the average model about the responses according to different environmental conditions, which is mostly related to air temperature, and they do not consider individual preferences. Thus, although the reactive and proactive control strategies act more intelligently by additionally considering the preferences of human occupants, we could not obtain explicit benefits to improve the satisfaction level and even in some cases, the solution quality may be harmed. On the other hand, the manual setting just make HVAC agents attempt to reach the desired temperature set point over time. Once HVAC agents get to the desired point, they are turned off, which will decrease the satisfaction level. If the tem-

[5]Note that the starting indoor temperature of the building is 55°F in the simulation, which causes the low average satisfaction level for a while.

perature is again away from the scope of desired temperature point, HVAC agents are turned on and the satisfaction level increases. As a result, the manual setting shows a race condition in the graph, which means it eventually cannot go over a certain point in terms of the satisfaction level. With revised equations considering more factors from the coordination perspective such as preferences, energy awareness, emotional contagion effect, etc., we expect more significant improvements in terms of the satisfaction level.

In our work, we still only separately consider two different optimization criteria — the energy consumption and the satisfaction level since this is still preliminary work. However, as we will eventually optimize multiple objectives in SAVES, we will be able to achieve effective multiagent team coordinations to minimize the total energy consumption while maximizing occupant's comfort level.

7. CONCLUSION

This paper aims to open a new area of research for multiagent systems: in many real-world problems, specifically in energy domains, we see many different levels of agent interactions and coordinations involved, and hence multiagent systems must address such complex situations to achieve the given objectives under uncertainty. In this work, we presented a new framework called SAVES based on distributed coordination reasoning for sustainability. There are three major new ideas in SAVES. SAVES: (i) explicitly considers uncertainty while reasoning about coordination in a distributed manner relying on MDPs; (ii) incorporates human behaviors and their occupancy preferences into planning and models them as part of the system; and (iii) evaluates various control strategies for multiagent teams on an existing university building as the practical research testbed with actual energy consumption data. We justified our design decisions in SAVES through a preliminary empirical evaluation and showed that SAVES can provide solutions to significantly reduce the energy consumption while achieving the comparable satisfaction level of building occupants. For future work, we will consider opportunities for direct occupant participation and incentivization via handheld devices and deploy our system to the real-world.

8. REFERENCES

[1] Method of testing for rating room air conditioner and packaged terminal air conditioner heating capacity. *ANSI/ASHRAE Standard 58*.

[2] Steering behaviors for autonomous characters. Website. `http://opensteer.sourceforge.net/`.

[3] Thermal environmental conditions for human occupancy. *ANSI/ASHRAE Standard 55*.

[4] Ventilation for acceptable indoor air quality. *ANSI/ASHRAE Standard 62*.

[5] S. Abras, S. Ploix, S. Pesty, and M. Jacomino. A multi–agent home automation system for power management. In *ICINCO*, 2006.

[6] V. Callaghan, G. Clarke, A. Pounds-Cornish, and S. Sharples. Buildings as intelligent autonomous systems : A model for integrating personal and building agents. In *International Conference on Intelligent Autonomous Systems (IAS)*, 2000.

[7] S. Caron and G. Kesidis. Incentive-based energy consumption scheduling algorithms for the smart grid. In *SmartGridComm*, 2010.

[8] W.-H. Chen and W.-S. Tseng. A novel multi-agent framework for the design of home automation. In *ITNG*, 2007.

[9] G. Conte and D. Scaradozzi. Viewing home automation systems as multiple agents systems. In *Multi-agent system for industrial and service robotics applications, RoboCUP*, 2003.

[10] D. Cook, M. Youngblood, and S. Das. A multi-agent approach to controlling a smart environment. *Designing Smart Homes*, 4008:165–182, 2006.

[11] D. J. Cook. Multi-agent smart environments. *Journal of Ambient Intelligence and Smart Environments*, 1:51–55, 2009.

[12] M. Dibley, H. Li, J. Miles, and Y. Rezgui. Towards intelligent agent based software for building related decision support. *Advanced Engineering Informatics*, 2010.

[13] J. Dijkstra and H. Timmermans. Towards a multi-agent model for visualizing simulated user behavior to support the assessment of design performance. *Automation in Construction*, 11(2):135 – 145, 2002.

[14] M. Fahrioglu and F. L. Alvardo. Designing incentive compatible contracts for effective demand managements. *IEEE Trans. Power Systems*, 15:1255–1260, 2000.

[15] Y. Gao, E. Tumwesigye, B. Cahill, and K. Menzel. Using data mining in optimization of building energy consumption and thermal comfort management. In *International Conference on Software Engineering and Data Mining (SEDM)*, 2010.

[16] Q. Khalid and R. Langhe. Evaluation and monitoring of energy consumption patterns using statistical modeling and simulation: Case study of commercial and residential buildings. In *International Conference on Emerging Technologies (ICET)*, 2010.

[17] K. Kim and K. J. Kim. Multi-agent-based simulation system for construction operations with congested flows. *Automation in Construction*, 19(7):867 – 874, 2010.

[18] A.-H. Mohsenian-Rad, V. Wong, J. Jatskevich, and R. Schober. Optimal and autonomous incentivebased energy consumption scheduling algorithm for smart grid. In *ISGT*, 2010.

[19] L. Perez-Lombard, J. Ortiz, and C. Pout. A review on buildings energy consumption information. *Energy and Buildings*, 40:394–398, 2008.

[20] Z. Ren, F. Yang, N. Bouchlaghem, and C. Anumba. Multi-disciplinary collaborative building design–a comparative study between multi-agent systems and multi-disciplinary optimisation approaches. *Automation in Construction*, 2011.

[21] A. H. Rosenfeld, J. J. Romm, H. Akbari, and A. C. Lloyd. Painting the town white – and green. *Technol. Rev.*, 100:52–59, 1997.

[22] N. Roy, A. Roy, and S. K. Das. Context-aware resource management in multi-inhabitant smart homes: A nash h-learning based approach. In *PerCom*, 2006.

[23] M. Santamouris, A. Argiriou, E. Dascalaki, C. Balaras, and A. Gaglia. Energy characteristics and savings potential in office buildings. *Solar Energy*, 52:59–66, 1994.

[24] P. Scerri, D. Pynadath, L. Johnson, P. Rosenbloom, M. Si, N. Schurr, and M. Tambe. A prototype infrastructure for distributed robot-agent-person teams. In *AAMAS*, 2003.

[25] H. S. Sun and S. E. Lee. Case study of data centers' energy performance. *Energy and Buildings*, 38:522–533, 2006.