A Preliminary Study on Preference Elicitation in DCOPs for Scheduling Devices in Smart Buildings*

Atena M. Tabakhi and **Ferdinando Fioretto** and **William Yeoh**
New Mexico State University, Las Cruces, NM 88003, USA
{amtabakh, ffiorett, wyeoh}@cs.nmsu.edu

Abstract

Distributed Constraint Optimization Problems (DCOPs) offer a powerful approach for the description and resolution of cooperative multi-agent problems. In such model a group of agents coordinate their actions to optimize a global objective function, taking into account their preferences and constraints. A core limitation of this model is the assumption that all agents' preferences are specified a priori. Unfortunately, in a number of application domains, such knowledge is not assumed, and these values may become available only after being elicited from users in the domain. Motivated by the current developments in smart buildings we explore the effect of preference elicitation in scheduling smart appliances within a network of interconnected buildings, with the goal of reducing the users' energy consumption costs, while taking into account the comfort of the occupants. This paper makes the following contributions: **(1)** It introduces the *Smart Building Devices Scheduling* (SBDS) problem and maps it as a DCOP; **(2)** It proposes a general model for preference elicitation in DCOPs; **(3)** and It empirically evaluates the effect of several heuristics to select a set of preferences to elicit in SBDS problems.

1 Introduction

The importance of constraint optimization is outlined by the impact of its application in a range of *Constraint Optimization Problems* (COPs), such as supply chain management [Rodrigues and Magatao, 2007] and roster scheduling [Abdennadher and Schlenker, 1999]. When resources are distributed among a set of autonomous agents and communication among the agents are restricted, COPs take the form of *Distributed Constraint Optimization Problems* (DCOPs) [Modi, 2003; Yeoh and Yokoo, 2012; Fioretto *et al.*, 2016a]. In this context, agents coordinate their value assignments to minimize the overall sum of resulting constraint costs. DCOPs are suitable to model problems that are distributed in nature and where a collection of agents attempts to optimize a global objective within the confines of localized communication. They have been employed to model various distributed optimization problems, such as meeting scheduling [Yeoh *et al.*, 2010; Zivan *et al.*, 2014], sensor networks [Farinelli *et al.*, 2008], coalition formation [Ueda *et al.*, 2010], and smart grids [Miller *et al.*, 2012].

The field of DCOP has matured significantly over the past decade since its inception [Modi *et al.*, 2005]. DCOP researchers have proposed a wide variety of solution approaches, from distributed search-based solvers [Modi *et al.*, 2005; Maheswaran *et al.*, 2004; Yeoh *et al.*, 2010] to distributed inference-based solvers [Petcu and Faltings, 2005; Vinyals *et al.*, 2011], as well as algorithms that use GPUs [Fioretto *et al.*, 2014] and logic programming [Le *et al.*, 2015; 2016] formulations. One of the core limitations of all these approaches is that they assume that the constraint costs in a DCOP are known a priori. Unfortunately, in some application domains, these costs are only known after they are queried or elicited from experts or users in the domain.

One such application is the smart device scheduling problem in a network of smart buildings, where the goal is to schedule a number of smart devices (e.g., smart thermostats, smart lightbulbs, smart washers, etc.) distributed across a network of smart buildings in such a way that optimizes the preferences of occupants in those buildings subject to a larger constraint that the peak energy demand in the network does not exceed a energy utility defined limit. We further describe this motivating application in more detail in Section 3.

DCOPs are a natural framework to represent this problem as each building can be represented as an agent and the preferences of occupants can be represented as constraints. Furthermore, due to privacy reasons, it is preferred that the preferences of each occupant are not revealed to other occupants. The DCOP formulation allows the preservation of such privacy since agents are only aware of constraints that they are involved in.

A priori knowledge on the constraint costs is unfeasible in our motivating application. A key challenge is thus in the elicitation of user preferences to populate the constraint cost tables. Due to the infeasibility of eliciting preferences to populate *all* preferences, in this paper, we introduce the *preference elicitation problem* for DCOPs, which studies *how to select a subset of k cost tables to elicit from each agent* with the goal of choosing those having a large impact on the overall solution quality. We propose several methods to select this subset of cost tables to elicit, based on the notion of *partial orderings*. Our preliminary results illustrate the effectiveness

*We would like to thank all the reviewers for their suggestions and helpful pointers. This research is partially supported by NSF grants 1345232, 1540168, and 1550662.

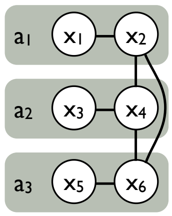

(a) Constraint Graph

for $i < j$

x_i	x_j	Costs
0	0	20
0	1	8
1	0	10
1	1	3

(b) Cost Table

Figure 1: Example DCOP.

of our approach in contrast to a baseline evaluator that randomly selects cost tables to elicit.

2 Background: Distributed Constraint Optimization Problems

A *Distributed Constraint Optimization Problem* (DCOP) is a tuple $\mathcal{P} = \langle \mathcal{X}, \mathcal{D}, \mathcal{F}, \mathcal{A}, \alpha \rangle$, where:

- $\mathcal{X} = \{x_1, \ldots, x_n\}$ is a set of *variables*;
- $\mathcal{D} = \{D_1, \ldots, D_n\}$ is a set of finite *domains* (i.e., D_i is the domain of x_i);
- $\mathcal{F} = \{f_1, \ldots, f_e\}$ is a set of *constraints* (also called *cost tables* in this work), where $f_i : \times_{x_j \in \mathbf{x}^{f_i}} D_i \rightarrow \mathbb{R}_0^+ \cup \{\bot\}$ maps each combination of value assignments of the variables $\mathbf{x}^{f_i} \subseteq \mathcal{X}$ in the *scope* of the function to a non-negative cost if the combination is allowed; $\bot$ is a special element used to denote that a given combination of value assignments is not allowed;
- $\mathcal{A} = \{a_1, \ldots, a_p\}$ is a set of *agents*;
- and $\alpha : \mathcal{X} \rightarrow \mathcal{A}$ is a function that maps each variable to one agent.

A *solution* σ is a value assignment to a set of variables $X_\sigma \subseteq \mathcal{X}$ that is consistent with the variables' domains. The cost function $\mathbf{F}_\mathcal{P}(\sigma) = \sum_{f \in \mathcal{F}, \mathbf{x}^f \subseteq X_\sigma} f(\sigma)$ is the sum of the costs of all the applicable constraints in σ. A solution is said to be *complete* if $X_\sigma = \mathcal{X}$ is the value assignment for *all* variables. The goal is to find an optimal complete solution $\mathbf{x}^* = \operatorname{argmin}_\mathbf{x} \mathbf{F}_\mathcal{P}(\mathbf{x})$.

Following Fioretto *et al.* [2016b], we introduce the following definitions:

Definition 1 *For each agent* $a_i \in \mathcal{A}$, $\mathbf{L}_i = \{x_j \in \mathcal{X} \mid \alpha(x_j) = a_i\}$ *is the set of its* local *variables.* $\mathbf{I}_i = \{x_j \in \mathbf{L}_i \mid \exists x_k \in \mathcal{X} \wedge \exists f_s \in \mathcal{F} : \alpha(x_k) \neq a_i \wedge \{x_j, x_k\} \subseteq \mathbf{x}^{f_s}\}$ *is the set of its* interface *variables.*

Definition 2 *For each agent* $a_i \in \mathcal{A}$*, its* local constraint graph $G_i = (\mathbf{L}_i, \mathcal{E}_{\mathcal{F}_i})$ *is a subgraph of the constraint graph, where* $\mathcal{F}_i = \{f_j \in \mathcal{F} \mid \mathbf{x}^{f_j} \subseteq \mathbf{L}_i\}$.

Figure 1(a) shows the constraint graph of a sample DCOP with 3 agents a_1, a_2, and a_3, where $\mathbf{L}_1 = \{x_1, x_2\}$, $\mathbf{L}_2 = \{x_3, x_4\}$, $\mathbf{L}_3 = \{x_5, x_6\}$, $\mathbf{I}_1 = \{x_2\}$, $\mathbf{I}_2 = \{x_4\}$, and $\mathbf{I}_3 = \{x_6\}$. The domains are $D_1 = \cdots = D_6 = \{0, 1\}$. Figure 1(b) shows the cost table of all constraints; all constraints have the same cost table for simplicity.

3 Scheduling of Devices in Smart Buildings

Through the proliferation of smart devices (e.g., smart thermostats, smart lightbulbs, smart washers, etc.) in our homes and offices, building automation within the larger smart grid is becoming inevitable. Building automation is the automated control of the building's devices with the objective of improved comfort of the occupants, improved energy efficiency, and reduced operational costs. In this paper, we are interested in *scheduling devices in smart buildings in a decentralized way*, where users are responsible for the schedule of the devices in their building, under the assumption that all the users cooperate to ensure that the total energy consumption of the neighborhood is within some limit defined by the energy provider such as a energy utility company.

We now provide a description of the *Smart Building Devices Scheduling* (SBDS) problem. We describe related solution approaches in Section 6. An SBDS problem is composed of a neighborhood $\mathcal{H}$ of smart buildings $h_i \in \mathcal{H}$ that are able to communicate with one another and whose energy demands are served by an energy provider. We assume that the provider sets energy prices according to a real-time pricing schema specified at regular intervals t within a finite time horizon H. We use $\mathbf{T} = \{1, \ldots, H\}$ to denote the set of time intervals and $\theta : \mathbf{T} \rightarrow \mathbb{R}^+$ to represent the price function associated with the pricing schema adopted, which expresses the cost per kWh of energy consumed by a consumer.

Within each smart building h_i, there is a set of (smart) electric devices $\mathcal{Z}_i$ networked together and controlled by a home automation system. All the devices are uninterruptible (i.e., they cannot be stopped once they are started). We use s_{z_j} and δ_{z_j} to denote, respectively, the start time and duration (expressed in multiples of time intervals) of device $z_j \in \mathcal{Z}_i$.

The energy consumption of each device z_j is ρ_{z_j} kWh for each hour that it is *on*. It will not consume any energy if it is *off*. We use the indicator function $\phi^t_{z_j}$ to indicate the state of the device z_j at time step t, and whose value is 1 exclusively when the device z_j is on at time step t:

$$\phi^t_{z_j} = \begin{cases} 1 & \text{if} \quad s_{z_j} \leq t \wedge s_{z_j} + \delta_{z_j} \geq t \\ 0 & \text{otherwise} \end{cases}$$

Additionally, the execution of a device z_j is characterized by a *cost* and a *discomfort value*. The cost represents the monetary expense for the user to schedule z_j at a given time, and we use C^t_i to denote the aggregated cost of the building h_i at time step t, expressed as:

$$C^t_i = P^t_i \cdot \theta(t), \tag{1}$$

where

$$P^t_i = \sum_{z_j \in \mathcal{Z}_i} \phi^t_{z_j} \cdot \rho_{z_j} \tag{2}$$

is the aggregate power consumed by building h_i at time step t. The discomfort value $\mu^t_{z_j} \in \mathbb{R}$ describes the degree of dissatisfaction for the user to schedule the device z_j at a given time step t. Additionally, we use U^t_i to denote the aggregated discomfort associated to the user in building h_i at time step t:

$$U^t_i = \sum_{z_j \in \mathcal{Z}_i} \phi^t_{z_j} \cdot \mu_{z_j}(t). \tag{3}$$

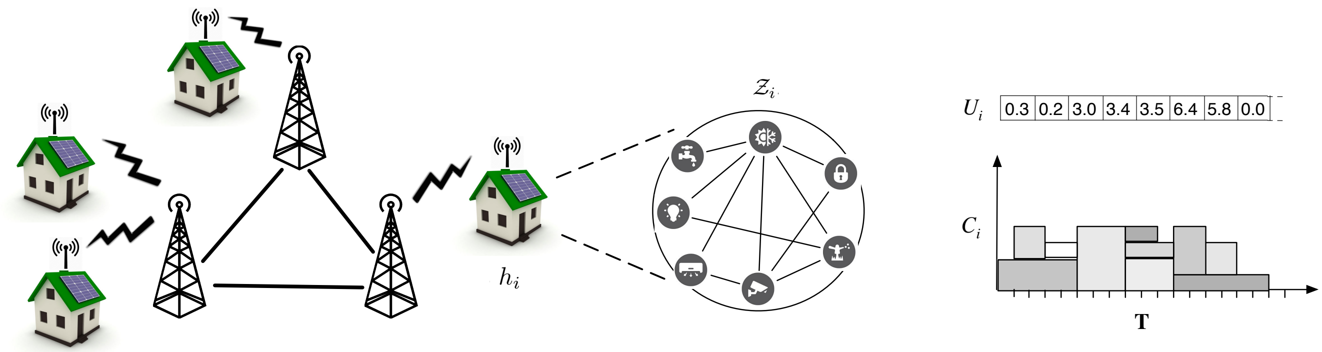

(a) (b) (c)

Figure 2: Smart Building Devices Scheduling illustration: (a) A network of smart buildings; (b) the set of smart devices $\mathcal{Z}_i$ controlled within a smart building h_i; and (c) an example of discomfort values (top) and costs (bottom) for a schedule of the devices in a building.

The SBDS problem is the problem of scheduling the devices of each building in the neighborhood in a coordinated fashion so as to minimize the monetary costs and, at the same time, minimize the discomfort of users. While this is a multi-objective optimization problem, we combine the two objectives into a single one through the use of a weighted sum:

$$\textbf{minimize} \sum_{t \in \mathbf{T}} \sum_{h_i \in \mathcal{H}} \alpha_c \cdot C_i^t + \alpha_u \cdot U_i^t \tag{4}$$

where α_c and α_u are weights in the open interval $(0, 1) \subseteq \mathbb{R}$ such that $\alpha_c + \alpha_u = 1$. The SBDS problem is also subject to the following constraints:

$$1 \leq s_{z_j} \leq T - \delta_{z_j} \qquad \forall h_i \in \mathcal{H}, z_j \in \mathcal{Z}_i \tag{5}$$

$$\sum_{t \in \mathbf{T}} \phi_{z_j}^t = \delta_{z_j} \qquad \forall h_i \in \mathcal{H}, z_j \in \mathcal{Z}_i \tag{6}$$

$$\sum_{h_i \in \mathcal{H}} P_i^t \leq \ell^t \qquad \forall t \in \mathbf{T} \tag{7}$$

where $\ell^t \in \mathbb{R}^+$ is the maximum allowed total energy consumed by all the buildings in the neighborhood at time step t. This constraint is typically imposed by the energy provider and is adopted to guarantee reliable electricity delivery. Constraint (5) expresses the lower and upper bounds for the start time associated to the schedule of each device. Constraint (6) ensures that the devices are scheduled and that they are executed for exactly their duration time. Constraint (7) represents the feasibility constraint and ensures that the total amount of energy consumed by the buildings in the neighborhood does not exceeds the maximum allowed threshold. Note that, in this work, we consider the devices to be *abstract* entities. Thus, scheduling a single device twice within the time horizon can be treated as two separate devices.

Figure 2 provides an illustration of the SBDS problem (a-b), and an example of the devices schedule for a building (c).

3.1 Preference Elicitation

While, in general, the real-time pricing schema θ that defines the cost per kWh of energy consumed and the energy consumption ρ_{z_j} of each device z_j are well-defined concepts and can be easily acquired or modeled, the preferences on the user's discomfort levels $\mu_{z_j}(t)$ on scheduling a device z_j at time step t are more subjective and, thus, more difficult to model explicitly.

We foresee two approaches to acquire these preferences: (1) eliciting them directly from the user and (2) estimating them based on historical preferences or from preferences of similar users. While the former method will be more accurate and reliable, it is cumbersome for the user to enter their preference for every device z_j at every time step t of the problem. Therefore, in this paper, we assume that a combination of the two approaches will be used, where a subset of preferences will be elicited and the remaining preferences will be estimated from historical sources or similar users.

4 DCOP Representation

We now describe how to map the SBDS problem to a DCOP:

- AGENTS: Each building $h_i \in \mathcal{H}$ is mapped to an agent $a_i \in \mathbf{A}$ in the DCOP.
- VARIABLES: For each building $h_i \in \mathcal{H}$, there are two types of variables for each agent:
  - The start time s_{z_j} of each device z_j is mapped to a decision variable x_j.
  - The indicator variables $\phi_{z_j}^t$ of each device z_j and time step t is mapped to an auxiliary variable x'_j.[1]
  - The aggregated energy consumed by all the devices in the building at each time step t is mapped to an auxiliary interface variable x''_j.

  All these variables are controlled by agent a_i and, thus, $\alpha(x_j) = \alpha(x'_j) = \alpha(x''_j) = a_i$ for all decision variables x_j and auxiliary variables x'_j and x''_j.
- DOMAINS: The domains of the decision and auxiliary variables are as follows:

[1] One can avoid to include such auxiliary variables by encoding the costs and discomfort values associated to the entire execution of a device in the problem constraints.

- The domain of decision variable x_j is the restricted set $\mathbf{T}$ as defined by Constraint (5).
- The domain of auxiliary variable x'_j is the set $\{0, 1\}$.
- The domain of auxiliary variable x''_j is the set $\{0, \ldots, \sum_{z_j \in \mathcal{Z}_i} \rho_{z_j}\}$.

- CONSTRAINTS: There are three types of constraints for each agent:
  - *Local* soft constraints (i.e., constraints that involve only variables controlled by the agent) whose costs correspond to the weighted summation of monetary costs and user discomfort, as defined by the objective function in Equation (4).
  - *Local* hard constraints that enforce Constraint (6).
  - *Global* hard constraints (i.e., constraints that involve variables controlled by different agents) that restrict the set of feasible aggregated energy consumption to be within the maximum allowed total energy consumed by all buildings, as defined by Constraint (7). Feasible aggregated consumptions incur a cost of 0 while infeasible aggregated consumptions incur a cost of $\perp$, which means that they are prohibited.

Since the DCOP enforces all Constraints (5), (6) and (7), an optimal complete DCOP solution that minimizes the sum of costs over all (local soft) constraints is *exactly* an optimal complete solution to the corresponding SBDS problem.

5 Preference Elicitation in DCOPs

As introduced in Section 3.1, one of the key drawbacks of existing DCOP approaches is that they assume that the cost tables of all constraints are known a priori, which is not the case for a number of real-world applications, including the SBDS problem. Due to the infeasibility of eliciting preferences to populate *all* cost tables, in this paper, we perform a preliminary study on how to choose a subset of k cost tables to populate. We first describe this optimization problem, and thus describe our proposed techniques.

Let $\hat{\mathcal{P}} = \langle \mathcal{X}, \mathcal{D}, \hat{\mathcal{F}}, \mathcal{A}, \alpha \rangle$ denote the DCOP whose constraints $\hat{\mathcal{F}}$ may have *inaccurate* cost tables. The constraints $\hat{\mathcal{F}} = \mathcal{F}_r \cup \mathcal{F}_u$ are composed of *revealed constraints* $\mathcal{F}_r$, whose cost tables are accurately revealed and reflect the actual user preferences, and *uncertain constraints* $\mathcal{F}_u$, whose cost tables are unrevealed and must be estimated from historical sources or similar users. All constraints that depend only on external parameters that are easily obtained are revealed constraints. We refer to this problem as the *uncertain DCOP*.

To abstract out the cost estimation in the uncertain constraints, we model those costs as random variables following a Normal distribution (e.g., one could fit a Normal distribution to the historical data).[2] Specifically, for each constraint $f \in \mathcal{F}_u$, the cost entry for the combination of values φ in its cost table is a random variable Y_φ obeying the Normal distribution $\mathcal{N}(\hat{\mu}_\varphi, \hat{\sigma}^2_\varphi)$ with mean $\hat{\mu}_\varphi$ and standard deviation $\hat{\sigma}_\varphi$.

Next, let $\mathcal{P} = \langle \mathcal{X}, \mathcal{D}, \mathcal{F}, \mathcal{A}, \alpha \rangle$ denote the DCOP whose constraints $\mathcal{F}$ have *accurate* cost tables, that is, they depend only on external parameters that are easily obtained (e.g., price function θ and energy consumption of devices ρ_{z_j}) or they depend on user preferences that are accurately obtained through an oracle. Here, we assume that the costs of each constraint $f \in \mathcal{F}$ that depend on user preferences are sampled from the same Normal distribution $\mathcal{N}(\hat{\mu}_\varphi, \hat{\sigma}^2_\varphi)$ in the corresponding uncertain DCOP. Therefore, the accurate costs are most likely the mean $\hat{\mu}_\varphi$ of the distribution but may have variations due to intricate subjective factors. We refer to this problem as the *oracle DCOP*.

[2] We leave the study of accurate cost estimation to future work.

5.1 Preference Elicitation Problem

The preference elicitation problem in DCOPs is formalized as follows: Given an oracle DCOP $\mathcal{P}$ and a value $k \in \mathbb{N}$, construct an uncertain DCOP $\hat{\mathcal{P}}$ that reveals only k constraints per agent (i.e., $|\mathcal{F}_r| = k \cdot |\mathbf{A}|$) and minimizes the error:

$$\epsilon_{\hat{\mathcal{P}}} = \mathbb{E}\left[|\mathbf{F}_{\hat{\mathcal{P}}}(\hat{\mathbf{x}}^*) - \mathbf{F}_{\mathcal{P}}(\mathbf{x}^*)|\right] \quad (8)$$

where $\hat{\mathbf{x}}^*$ is the optimal solution for a *realization* of the uncertain DCOP $\hat{\mathcal{P}}$, and $\mathbf{x}^*$ is the optimal complete solution for the oracle DCOP $\mathcal{P}$. A realization of an uncertain DCOP $\hat{\mathcal{P}}$ is a DCOP (with no uncertainty), whose costs of each combinations of values φ in each uncertain constraint, are realization of the random variables Y_φ of $\hat{\mathcal{P}}$.

5.2 Preference Elicitation Heuristics

Note that the possible numbers of uncertain DCOPs that can be generated is $\binom{|\mathcal{F}|}{k \cdot |\mathbf{A}|}$. Since solving each DCOP is NP-hard [Modi, 2003], the preference elicitation problem is a particularly challenging one. Thus, we propose five heuristic methods to determine the subset of functions to reveal, so to construct an uncertain problem $\hat{\mathcal{P}}$.

Let us first introduce a general concept of partial ordering between cost tables of uncertain constraints.

Definition 3 *Given a* partial ordering $\circ$ *on the uncertain set* $\mathcal{F}_u \subset \hat{\mathcal{F}}$*, and two cost tables of uncertain constraints* $f_i, f_j \in \mathcal{F}_u$*, we say that* f_i dominates f_j *according to* $\circ$ *if* $f_i \succeq_\circ f_j$.

We now introduce the heuristic methods to choose the first k uncertain constraints ordered by the relation $\succeq_\circ$.

Average of the Expected Costs

If the ordering $\circ = \mathbb{EE}[\cdot]$ is done according to the average of the expected costs of the uncertain constraints, then, given two unknown functions $f_i, f_j \in \mathcal{F}_u$, we say that $f_i \succeq_{\mathbb{EE}} f_j$ *iff*:

$$\mathbb{EE}[f_i] \leq \mathbb{EE}[f_j]$$

where

$$\mathbb{EE}[f_i] = \frac{1}{|\Sigma^{f_i}_{\mathbf{x}}|} \sum_{\varphi \in \Sigma^{f_i}_{\mathbf{x}}} \hat{\mu}_\varphi$$

and $\Sigma^{f_i}_{\mathbf{x}}$ is the set of all the possible value assignments for the variables in $\mathbf{x}^{f_i}$.

Average of the Variance

If the ordering $\circ = \mathbb{E}\pi[\cdot]$ is done according to the average variance of the uncertain constraints, then, given two unknown functions $f_i, f_j \in \mathcal{F}_u$, we say that $f_i \succeq_{\mathbb{E}\pi} f_j$ *iff*:

$$\mathbb{E}\pi[f_i] \leq \mathbb{E}\pi[f_j]$$

where

$$\mathbb{E}\pi[f_i] = \frac{1}{|\Sigma_{\mathbf{x}}^{f_i}|} \sum_{\varphi \in \Sigma_{\mathbf{x}}^{f_i}} \hat{\sigma}_{\varphi}^2.$$

Variance of the Expected Costs
If the ordering $\circ = \pi\mathbb{E}[\cdot]$ is done according to the variance of expected costs of the uncertain constraints, then, given two unknown functions $f_i, f_j \in \mathcal{F}_u$, we say that $f_i \succeq_{\pi\mathbb{E}} f_j$ *iff*:

$$\pi\mathbb{E}[f_i] \leq \pi\mathbb{E}[f_j]$$

where

$$\pi\mathbb{E}[f_i] = \frac{1}{|\Sigma_{\mathbf{x}}^{f_i}|} \sum_{\varphi \in \Sigma_{\mathbf{x}}^{f_i}} (\hat{\mu}_{\varphi} - \mathbb{E}\mathbb{E}[f_i])^2$$

Variance of the Variance
If the ordering $\circ = \pi\pi[\cdot]$ is done according to the variance of variance of the uncertain constraints, then, given two unknown functions $f_i, f_j \in \mathcal{F}_u$, we say that $f_i \succeq_{\pi\pi} f_j$ *iff*:

$$\pi\pi[f_i] \leq \pi\pi[f_j]$$

where

$$\pi\pi[f_i] = \frac{1}{|\Sigma_{\mathbf{x}}^{f_i}|} \sum_{\varphi \in \Sigma_{\mathbf{x}}^{f_i}} \left(\hat{\sigma}_{\varphi}^2 - \mathbb{E}\pi[f_i]\right)^2$$

Second-Order Stochastic Dominance
In general, it is well-known (particularly in financial domains) that maximizing cost without considering risk does not yield good solutions in risky environments. To incorporate the notion of risk while ordering the uncertain constraints, we use the concept of second-order stochastic dominance [Levy, 1998].

If the ordering $\circ = SD[\cdot]$ is done according to the stochastic dominance criteria, then, given two unknown functions $f_i, f_j \in \mathcal{F}_u$, we say that $f_i \succeq_{SD} f_j$ *iff*:

$$\sum_{m=1}^{x} (f_i(m) - f_j(m)) \geq 0$$

for all values of $x \leq |\Sigma_{\mathbf{x}}^{f_i}|$, where $f_i(m) = \hat{\mu}_t$ is the expected cost of the m-th value assignment for the variables in $\mathbf{x}^{f_i}$.

Notice however, that both functions may not stochastically dominate each other. Furthermore, this ordering is only defined if the number of possible value assignments for both unknown functions are identical, i.e., $|\Sigma_{\mathbf{x}}^{f_i}| = |\Sigma_{\mathbf{x}}^{f_j}|$. In our DCOP model of the SBDS problem, this is the case if the horizon to schedule all devices are identical.

6 Related Work

The problem of scheduling devices in smart buildings has recently attracted large interest within the AI and the smart grid communities. Georgievski *et al.* [2012] proposed a system to monitor and control electrical appliances in a building with the objective of reducing the energy bill costs. Scott *et al.* [2013] have also studied a centralized online stochastic optimization approach for a home automation system as a demand response mechanism, where the uncertainty comes from future prices, occupant behavior, and environmental conditions. Sou *et al.* [2011] proposed a Mixed Integer Linear Program (MILP) to address smart appliance scheduling problem using low granularity for the technical specification of smart appliance (e.g., they distinguish in the various energy phases carried in a dishwasher, or washing machine cycle). However, due to the high complexity of the problem they suggest to adopt suboptimal solutions to reduce the overall solving time. Another proposal to speed up the resolution time of a MILP formulation for scheduling smart devices was presented by Tsui and Chan [2012]. The authors study a relaxation of a MILP formulation for the automatic load management of appliances in a smart home as a convex program optimization, which speeds up the resolution process, however they provide no guarantees on the solution quality with respect to the original problem. Unlike our approach, these proposals focus on single building problems, and/or are inherently centralized.

Scott and Thiébaux [2015] have also studied a distributed demand response mechanism for the scheduling of shiftable loads in smart homes, within a non-convex optimization context, and show that simple approaches—even if not guaranteed to find feasible solutions—can be effective. However, they do not address the occupant's comfort while computing the appliances policies and focus exclusively on load management.

The problem of preference elicitation in DCOPs is related to a class of DCOPs where agents have partial knowledge on the costs of their constraints and, therefore, they may discover the unknown costs via exploration [Taylor *et al.*, 2011; Zivan *et al.*, 2015]. In this context, agents must balance the coordinated *exploration* of the unknown environment and the *exploitation* of the known portion of the rewards, in order to optimize the global objective [Stranders *et al.*, 2012]. Another orthogonal related DCOP model is the problem where costs are sampled from probability distribution functions [Nguyen *et al.*, 2014]. In such a problem, agents seek to minimize either the worst-case regret [Wu and Jennings, 2014] or the expected regret [Le *et al.*, 2016].

7 Empirical Evaluation

We evaluate the effect of preference elicitation in DCOPs in synthetic SBDS problems. In our experiment we consider $|\mathcal{H}| = 10$ buildings, each controlling $|\mathcal{Z}_i| = 10$ smart devices. The list of smart devices adopted (i.e., $\mathcal{Z} = \cup_i \mathcal{Z}_i$) is illustrated in Table 1, including their power consumption ρ_{z_j} (in kWh) and duration δ_j (in minutes) for each device $z_j \in \mathcal{Z}$. The device specifics (i.e., duration and power consumption) follow those by Zhang *et al.* [2013].

We populate the set of smart devices $\mathcal{Z}_i$ of each building by randomly sampling 10 elements from $\mathcal{Z}$. Thus, a building might control multiple devices of the same type. In our experiment, we set a time horizon $H = 12$ with increments of 30 minutes. The values for the real time pricing schema $\theta(t)$ are summarized in Table 2.

For each building and each device z_j, the user preferences on the discomfort values $\mu_{z_j}^t$ are generated from a Normal distribution $\mathcal{N}(\hat{\mu}, \hat{\sigma}^2)$, with $\hat{\mu}$ randomly sampled in $[1, 100]$, and $\hat{\sigma}^2$ in $[1, \frac{\sqrt{\hat{\mu}}}{2}]$. Finally, the weights α_c and α_u of the ob-

Device	Power (kWh)	Duration (min.)
dish washer	0.75	120
washing machine	1.20	90
dryer	2.50	60
cooker hob	3.0	30
cooker oven	5.0	30
microwave	1.70	30
laptop	0.10	120
desktop computer	0.30	180
vacuum cleaner	1.20	30
fridge	0.30	360
electrical vehicle	3.50	180

Table 1: Smart devices with their power consumption and schedule duration.

Time (min.)	[0-60]	[60-120]	[120-180]	[180-240]	[240-300]	[300-360]
RTP ($/kWh)	0.172	0.161	0.191	0.145	0.149	0.174

Table 2: Real time pricing schema.

jective function defined in Equation (4) are set to 0.5. These settings are employed to create both an oracle DCOP and the corresponding uncertain DCOP, as described in Section 5, except that the values of the constraints of the uncertain DCOPs are not realized (i.e., they are distributions). All the problems are modeled and solved optimally using MiniZinc [Nethercote *et al.*, 2007] on an *Intel Core i7-3770* CPU 3.40GHz with16 GB of RAM.

Figure 3 illustrates the results on the error corresponding to the preference elicitation problem for various number k of constraints to elicit per agent, and with respect to the partial orderings: Average of the Expected Costs (AE), Average of the Variance (AV), Variance of the Expected Costs (VE), Variance of the Variance (VV), and Second-Order Stochastic Dominance (SD), as described in Section 5.2. Additionally, we employ a Random (RN) heuristic, as baseline for comparison, which chooses the k constraints to elicit per agent at random. We report the normalized error $\frac{\epsilon_{\hat{\mathcal{P}}}}{\mathbf{F}_{\mathcal{P}}(\mathbf{x}^*)}$, where $\epsilon_{\hat{\mathcal{P}}}$ is the error as defined by Equation (8). An accurate computation of this error requires us to generate all possible realizations for the uncertain DCOPs. Due to the complexity of such task, we create $m = 50$ realizations of the uncertain DCOPs and compute the error $\epsilon_{\hat{\mathcal{P}}}$ in this reduced sampled space. We control k so that the percentage of the constraints elicited from the oracle DCOP ($\frac{|\mathcal{F}_e|}{|\mathcal{F}|} \cdot 100$) varies from 20% to 80%. For each setting, the values of the uncertain constraints $\mathcal{F}_u$ of the resulting uncertain DCOP obey the same distributions used to model the corresponding realizations in the oracle DCOP. The results are averaged over 50 randomly generated SBDS problem instances. We make the following observations:

- As expected, for all the partial orderings tested, the error decreases as the number of cost tables to elicit increases.
- AE and AV outperform all other heuristics.

8 Conclusions

Motivated by the current developments in smart buildings, we explore the effect of preference elicitation in scheduling

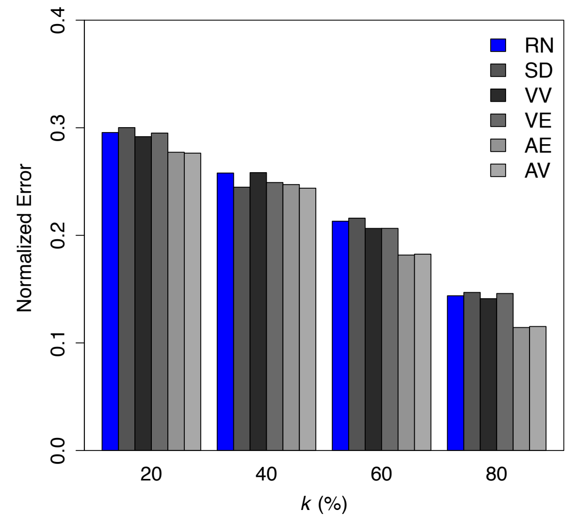

Figure 3: Preference Elicitation Results.

smart appliances within a network of interconnected buildings, with the goal of reducing the users' energy consumption costs, while taking into account the comfort of the occupants. After modeling this problem as a DCOP, we propose a general model for preference elicitation in DCOPs. Due to the infeasibility of eliciting preferences to populate *all* DCOP cost tables, we proposed several methods to select a subset of k cost tables to elicit per agent, based on the notion of partial orderings. Our preliminary results show that our best methods are more accurate than a baseline method that randomly selects cost tables to elicit.

Future work will focus on an extensive analysis of the proposed methods on a more realistic setting for the SBDS agents as well as incorporating state-of-the-art methods for predicting energy consumption in homes [Truong *et al.*, 2013; Lachut *et al.*, 2014].

References

[Abdennadher and Schlenker, 1999] S. Abdennadher and H. Schlenker. Nurse Scheduling using Constraint Logic Programming. In *Proceedings of the Conference on Innovative Applications of Artificial Intelligence (IAAI)*, pages 838–843, 1999.

[Farinelli *et al.*, 2008] A. Farinelli, A. Rogers, A. Petcu, and N. Jennings. Decentralised coordination of low-power embedded devices using the Max-Sum algorithm. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 639–646, 2008.

[Fioretto *et al.*, 2014] F. Fioretto, T. Le, W. Yeoh, E. Pontelli, and T. C. Son. Improving DPOP with branch consistency for solving distributed constraint optimization problems. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*, pages 307–323, 2014.

[Fioretto *et al.*, 2016a] F. Fioretto, E. Pontelli, and W. Yeoh. Distributed constraint optimization problems and applications: A survey. *CoRR*, abs/1602.06347, 2016.

[Fioretto *et al.*, 2016b] F. Fioretto, W. Yeoh, and E. Pontelli. Multi-variable agent decomposition for dcops. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2016.

[Georgievski *et al.*, 2012] I. Georgievski, V. Degeler, G. A. Pagani, T. A. Nguyen, A. Lazovik, and M. Aiello. Optimizing energy costs for offices connected to the smart grid. *IEEE Transactions on Smart Grid*, 3(4):2273–2285, 2012.

[Lachut *et al.*, 2014] D. Lachut, N. Banerjee, and S. Rollins. Predictability of energy use in homes. In *Proceedings of the International Green Computing Conference (IGCC)*, pages 1–10, 2014.

[Le *et al.*, 2015] T. Le, T. C. Son, E. Pontelli, and W. Yeoh. Solving distributed constraint optimization problems with logic programming. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2015.

[Le *et al.*, 2016] T. Le, F. Fioretto, W. Yeoh, T. C. Son, and E. Pontelli. ER-DCOPs: A framework for distributed constraint optimization with uncertainty in constraint utilities. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2016.

[Levy, 1998] H. Levy. *Stochastic Dominance – Investment Decision Making under Uncertainty (Studies in Risk and Uncertainty)*. Springer, 1998.

[Maheswaran *et al.*, 2004] R. Maheswaran, J. Pearce, and M. Tambe. Distributed algorithms for DCOP: A graphical game-based approach. In *Proceedings of the International Conference on Parallel and Distributed Computing Systems (PDCS)*, pages 432–439, 2004.

[Miller *et al.*, 2012] S. Miller, S. Ramchurn, and A. Rogers. Optimal decentralised dispatch of embedded generation in the smart grid. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 281–288, 2012.

[Modi *et al.*, 2005] P. Modi, W.-M. Shen, M. Tambe, and M. Yokoo. ADOPT: Asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence*, 161(1–2):149–180, 2005.

[Modi, 2003] P. Modi. *Distributed Constraint Optimization for Multiagent Systems*. PhD thesis, University of Southern California, Los Angeles (United States), 2003.

[Nethercote *et al.*, 2007] N. Nethercote, P. Stuckey, R. Becket, S. Brand, G. Duck, and G. Tack. Minizinc: Towards a standard cp modelling language. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*, pages 529–543, 2007.

[Nguyen *et al.*, 2014] D. T. Nguyen, W. Yeoh, H. C. Lau, S. Zilberstein, and C. Zhang. Decentralized multi-agent reinforcement learning in average-reward dynamic DCOPs. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 1447–1455, 2014.

[Petcu and Faltings, 2005] A. Petcu and B. Faltings. A scalable method for multiagent constraint optimization. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1413–1420, 2005.

[Rodrigues and Magatao, 2007] L.C.A. Rodrigues and L. Magatao. Enhancing Supply Chain Decisions Using Constraint Programming: A Case Study. In *Proceedings of the Mexican International Conference on Artificial Intelligence (MICAI)*, pages 1110–1121, 2007.

[Scott and Thiébaux, 2015] P. Scott and W. Thiébaux. Distributed multi-period optimal power flow for demand response in microgrids. In *Proceedings of the International Conference on Future Energy Systems (e-Energy)*, pages 17–26, 2015.

[Scott *et al.*, 2013] P. Scott, S. Thiébaux, M. van den Briel, and P. van Hentenryck. Residential demand response under uncertainty. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*, pages 645–660, 2013.

[Sou *et al.*, 2011] K. C. Sou, J. Weimer, H. Sandberg, and K. H. Johansson. Scheduling smart home appliances using mixed integer linear programming. In *IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, pages 5144–5149. IEEE, 2011.

[Stranders *et al.*, 2012] R. Stranders, F. Delle Fave, A. Rogers, and N. Jennings. DCOPs and bandits: Exploration and exploitation in decentralised coordination. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 289–297, 2012.

[Taylor *et al.*, 2011] M. Taylor, M. Jain, P. Tandon, M. Yokoo, and M. Tambe. Distributed on-line multi-agent optimization under uncertainty: Balancing exploration and exploitation. *Advances in Complex Systems*, 14(03):471–528, 2011.

[Truong *et al.*, 2013] N. C. Truong, J. McInerney, L. Tran-Thanh, E. Costanza, and S. Ramchurn. Forecasting multi-appliance usage for smart home energy management. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2908–2914, 2013.

[Tsui and Chan, 2012] K. M. Tsui and S.-C. Chan. Demand response optimization for smart home scheduling under real-time pricing. *IEEE Transactions on Smart Grid*, 3(4):1812–1821, 2012.

[Ueda *et al.*, 2010] S. Ueda, A. Iwasaki, and M. Yokoo. Coalition structure generation based on distributed constraint optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 197–203, 2010.

[Vinyals *et al.*, 2011] M. Vinyals, J. Rodríguez-Aguilar, and J. Cerquides. Constructing a unifying theory of dynamic programming DCOP algorithms via the generalized distributive law. *Autonomous Agents and Multi-Agent Systems*, 22(3):439–464, 2011.

[Wu and Jennings, 2014] F. Wu and N. Jennings. Regret-based multi-agent coordination with uncertain task rewards. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 1492–1499, 2014.

[Yeoh and Yokoo, 2012] W. Yeoh and M. Yokoo. Distributed problem solving. *AI Magazine*, 33(3):53–65, 2012.

[Yeoh *et al.*, 2010] W. Yeoh, A. Felner, and S. Koenig. BnB-ADOPT: An asynchronous branch-and-bound DCOP algorithm. *Journal of Artificial Intelligence Research*, 38:85–133, 2010.

[Zhang *et al.*, 2013] D. Zhang, N. Shah, and L. Papageorgiou. Efficient energy consumption and operation management in a smart building with microgrid. *Energy Conversion and Management*, 74:209–222, 2013.

[Zivan *et al.*, 2014] R. Zivan, S. Okamoto, and H. Peled. Explorative anytime local search for distributed constraint optimization. *Artificial Intelligence*, 212:1–26, 2014.

[Zivan *et al.*, 2015] R. Zivan, H. Yedidsion, S. Okamoto, R. Glinton, and K. Sycara. Distributed constraint optimization for teams of mobile sensing agents. *Journal of Autonomous Agents and Multi-Agent Systems*, 29(3):495–536, 2015.