Preference Elicitation in DCOPs for Scheduling Devices in Smart Buildings

Atena M. Tabakhi
Department of Computer Science
New Mexico State University
amtabakh@cs.nmsu.edu

Abstract

Researchers have used *Distributed Constraint Optimization Problems* (DCOPs) as a powerful approach to model various multi-agent coordination problems, taking into account their preferences and constraints. A core limitation of this model is the assumption that all agents' preferences are specified a priori. However, in a number of application domains such knowledge become available only after being elicited from users in these domains. In this abstract, we explore the effects of preference elicitation in our motivating application of scheduling smart appliances with the aim of reducing users' electricity bill cost as well as increasing their comfort.

Introduction

Distributed Constraint Optimization Problems (DCOPs) (Modi 2003; Yeoh and Yokoo 2012) are suitable to model problems that are distributed in nature, such as nurse scheduling, meeting scheduling, and supply chain management. While the field has matured significantly over the past decade, existing DCOPs assume that their constraint utilities are known a priori. In some application domains, these utilities are only known after they are queried or elicited from experts or users in the domain.

Our motivating application is the smart device scheduling problem in a network of smart buildings. The goal is to schedule a number of smart devices such as smart thermostat and smart washers that optimizes the occupants' preferences of those building subject to that the peak energy demand does not exceed the energy utility defined limit.

A priori knowledge on the constraint utilities is not feasible in our motivating application. Therefore, the key challenge is in the elicitation of user preferences to populate the constraints cost tables. We propose several methods to select a subset of k cost tables to elicit, from each agent, with the highest impact on the solution quality.

Scheduling of Devices in Smart Buildings

The Smart Building Devices Scheduling (SBDS) problem is composed of a neighborhood $\mathcal{H}$ of smart buildings that are able to communicate with one another, whose energy demands are served by an energy provider. We use $\mathbf{T} = \{1, \dots, H\}$ to denote the set of time intervals and $\theta : \mathbf{T} \to \mathbb{R}^+$ to represent the price function associated with the pricing schema defined by an energy provider. Within each smart building $h_i \in \mathcal{H}$, 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 multiple 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, whose value is 1 exclusively when the device z_j is on at time step t.

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. We use C^t_i to denote the aggregated cost of the building h_i at time step t, expressed as: $C^t_i = P^t_i \cdot \theta(t)$, where $P^t_i = \sum_{z_j \in \mathcal{Z}_i} \phi^t_{z_j} \cdot \rho_{z_j}$ 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: $U^t_i = \sum_{z_j \in \mathcal{Z}_i} \phi^t_{z_j} \cdot \mu_{z_j}(t)$.

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 objective through the use of a weighted sum:

$$\textbf{minimize} \sum_{t \in \mathbf{T}} \sum_{h_i \in \mathcal{H}} \alpha_c \cdot C^t_i + \alpha_u \cdot U^t_i, \tag{1}$$

where α_c and α_u are weights in the open interval $(0, 1) \subseteq \mathbb{R}$ such that $\alpha_c + \alpha_u = 1$.

DCOP Representation

We map the SBDS problem to a DCOP so that each building in the SBDS problem is mapped to an agent in our DCOP model. The start time of each device is mapped to a decision variable. The domains of decision variables are restricted set $\mathbf{T}$. There are three types of constraints for each agent in our DCOP model: local soft constraints, local hard constraints, and global hard constraints. For simplicity we only explain the local soft constraints. The constraints that involve only variables, controlled by the agent, whose costs correspond to the weighted summation of monetary costs and user discomfort are local soft constraints. An optimal complete DCOP solution that minimizes the sum of costs over all local soft constraints is exactly an optimal solution to the corresponding SBDS problem.

Preference Elicitation in DCOPs

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. We propose the techniques of choosing a subset of k cost tables to populate, due to the infeasibility of eliciting preferences to populate *all* cost tables.

We introduce the uncertain DCOP $\hat{\mathcal{P}}$ whose constraints $\hat{\mathcal{F}}$ may have *inaccurate* cost tables. *Inaccurate* cost tables are composed of *revealed constraints* $\mathcal{F}_r$, with cost tables that reflect the actual user preferences, and *uncertain constraints* $\mathcal{F}_u$, with cost tables that must be estimated from historical sources or similar users. Next, we present the oracle DCOP $\mathcal{P}$ whose constraints $\mathcal{F}$ have *accurate* cost tables that depend on accurately obtained user preferences. We assume that the costs of each constraint $f \in \mathcal{F}$, which 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.

Preference Elicitation Problem

The preference elicitation problem in DCOPs is formalized as follows: Given an oracle DCOP $\mathcal{P}$ and a value $k \in \mathbb{N}$, we 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:

$$\epsilon_{\hat{\mathcal{P}}} = \mathbb{E}\left[|\mathbf{F}_{\mathcal{P}}(\hat{\mathbf{x}}^*) - \mathbf{F}_{\mathcal{P}}(\mathbf{x}^*)|\right]. \quad (2)$$

In this equation, $\hat{\mathbf{x}}^*$ 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 combination of values φ in each uncertain constraint, are realization of the random variables Y_φ of $\hat{\mathcal{P}}$.

We define a general concept of partial ordering between cost tables of uncertain constraints. 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 $f_i \succeq_\circ f_j$. What follows are the proposed heuristic methods to choose the first k uncertain constraints ordered by the relation $\succeq_\circ$.

- *Average of the Variance* (AV): we say that $f_i \succeq_{\mathbb{E}\pi} f_j$ *iff*: $\mathbb{E}\pi[f_i] \leq \mathbb{E}\pi[f_j]$, where $\mathbb{E}\pi[f_i] = \frac{1}{|\Sigma^{f_i}_{\mathbf{x}}|}\sum_{\varphi \in \Sigma^{f_i}_{\mathbf{x}}} \hat{\sigma}^2_\varphi$ and $\Sigma^{f_i}_{\mathbf{x}}$ is the set of all the possible value assignments for the variables in $\mathbf{x}^{f_i}$.

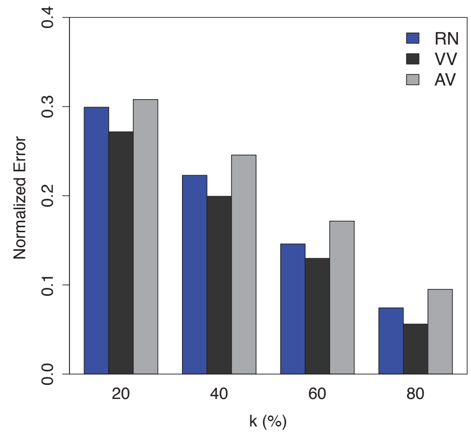

Figure 1: Preference Elicitation Results.

- *Variance of the Variance* (VV): we say that $f_i \succeq_{\pi\pi} f_j$ *iff*: $\pi\pi[f_i] \leq \pi\pi[f_j]$, where $\pi\pi[f_i] = \frac{1}{|\Sigma^{f_i}_{\mathbf{x}}|}\sum_{\varphi \in \Sigma^{f_i}_{\mathbf{x}}} \left(\hat{\sigma}^2_\varphi - \mathbb{E}\pi[f_i]\right)^2$

Empirical Evaluation

We evaluate the effect of preference elicitation in DCOPs in the synthetic SBDS problem. In our experiment we consider $|\mathcal{H}| = 10$ buildings, each controlling $|\mathcal{Z}_i| = 10$ smart devices. We set a time horizon $H = 12$ with increments of 30 minutes. Finally, the weights α_c and α_u of the objective function defined in Equation (1) are set to 0.5. 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 (2).

Figure 1 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. Additionally, we employ a Random (RN) heuristic, as baseline for comparison, which chooses the k constraints to elicit per agent at random. 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.

Conclusions

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.

References

Modi, P. 2003. *Distributed Constraint Optimization for Multiagent Systems*. Ph.D. Dissertation, University of Southern California, Los Angeles (United States).

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