

A Distributed Constraint Optimization Approach for Coordination under Uncertainty

(Extended Abstract)

James Atlas
Computer and Information Sciences
University of Delaware
Newark, DE 19716
atlas@cis.udel.edu

ABSTRACT

Distributed Constraint Optimization (DCOP) provides a rich framework for modeling multi-agent coordination problems. We present a mapping to DCOP for large-scale team coordination problems that were used in the DARPA Coordinators program. We develop a new DCOP algorithm that scales to this domain and show that our algorithm outperforms other DCOP algorithms for this domain. Our approach is competitive with other general approaches used in the DARPA Coordinators program.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—Multiagent Systems

General Terms

Algorithms, Experimentation

Keywords

Distributed Constraint Optimization, Multi-agent Coordination, Task Scheduling

1. INTRODUCTION

Distributed Constraint Optimization (DCOP) is a general problem representation for multi-agent systems. Recent advances in DCOP algorithm development have led to an increasing number of application domains and focus on DCOP techniques. Recent applications of DCOP to real-world problems include sensor networks[2], traffic flow cooperation [4], and meeting scheduling[5]. These existing problem domains for DCOP focus on small (<100 variables), deterministic domains. We present a mapping to DCOP for large-scale team coordination problems that were used in the DARPA Coordinators program, an algorithm to solve this mapping, and comparison of the results with other DCOP algorithms and DARPA team approaches.

Cite as: A Distributed Constraint Optimization Approach for Coordination under Uncertainty, (Extended Abstract), James Atlas, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. 1263–1264
Copyright © 2009, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

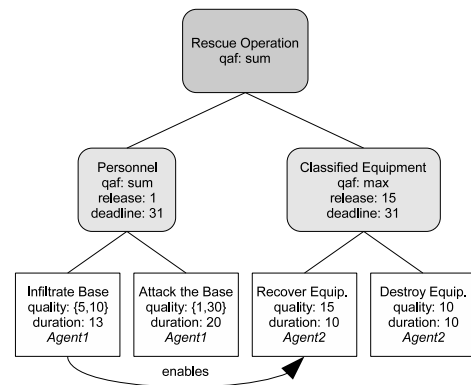


Figure 1: An example C-TÆMS problem instance.

2. C-TÆMS COORDINATION PROBLEM

Multi-agent task planning and scheduling problems require a rich language for domain representation. The original TÆMS (Task Analysis, Environment Modeling, and Simulation) language was developed to provide a domain independent, quantitative representation of the complex coordination problem [3]. A C-TÆMS problem instance contains a set of agents and a hierarchically decomposed task structure. Nodes in the graph are either complex tasks (internal nodes) or primitive methods (leaf nodes). Each node may have temporal constraints on the earliest start time and the deadline. Nodes may also have non-local effect (NLE) constraints that represent hard (enables and disables) and soft (facilitates and hinders) node relationships. Methods have probabilistic outcomes for duration, quality, and cost. A sample C-TÆMS problem instance is shown in Figure 1.

3. MAPPING C-TÆMS TO DCOP

DCOP has been formalized in slightly different ways in recent literature, so we will adopt the definition as presented in [5]. A Distributed Constraint Optimization Problem with n nodes and m constraints consists of the tuple $\langle X, D, U \rangle$ where:

- $X = \{x_1, \dots, x_n\}$ is a set of variables, each one assigned to a unique agent
- $D = \{d_1, \dots, d_n\}$ is a set of finite domains for each variable

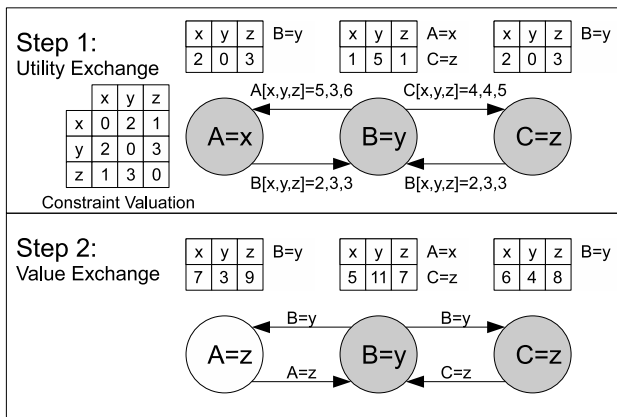


Figure 2: DNEA in action. Given constraint valuations shown at top and a random starting assignment of $A=x$, $B=y$, and $C=z$, DNEA finds the optimal assignment in one round of exchanges.

- $U = \{u_1, \dots, u_m\}$ is a set of utility functions such that each function involves a subset of variables in X and defines a utility for each combination of values among these variables (constraints)

An optimal solution to a DCOP instance consists of an assignment of values in D to X such that the sum of utilities in U is maximal. Our mapping for C-TÆMS to DCOP contains two parts: variable and constraint mappings.

Variables are created for each method and task in the C-TÆMS problem. Method variables are created with all possible start times as values and a value for *not scheduled*. Task variables can have several different sets of values depending on the type of QAF assigned to the task; these reflect whether a task is enabled for execution, if it accumulates quality, and if it is being forced to execute/not execute. In addition, a special end-time variable is created for each task with an outgoing NLE at or above it in the structure.

Constraints are created between each related node in the problem structure. Task-subtask and task-method constraints define how quality accumulation contributes to overall solution utility. Method-method mutex constraints enforce that agents only perform one method at a time. NLE constraints between nodes enforce hard (enables and disables) and soft (facilitates and hinders) node relationships.

4. DISTRIBUTED NEIGHBOR EXCHANGE

Our Distributed Neighbor Exchange Algorithm (DNEA) is similar in phases to other local value exchange based DCOP algorithms, including MGM, DSA/SCA, DBA, and max-sum [1, 2]. These algorithms exchange current variable assignments with neighbors, compute a maximization function based on neighboring assignments, and then choose to update or not update the local variable assignment. DNEA’s main advantage is that it scales to large problems, operating in $O(\sum_{Y \in N_X} |D_Y| \cdot |D_X|)$ per cycle, while finding high-utility solutions using approximate neighborhood optimization. An example is shown in Figure 2.




Figure 3: Estimated Solution Utility as a sum of DCOP utility for static scenario execution after 100 cycles, with and without an initial schedule. Closer to zero is better.




Figure 4: OptOP5PMix: Solution Quality as % of optimal (as determined by an offline solver).

5. RESULTS

The DARPA COORDINATORs project uses a simulation framework based on simulated time ticks with dynamic execution outcomes. In Figure 3 we show comparison results between DNEA and other scalable local search algorithms (MGM2 and DSA2) using our C-TÆMS mapping in a static scenario. Results of our full solution on a subset of the real DARPA test scenarios are shown in Figure 4. We show comparisons with a naive approach that follows a static schedule, and with the three contracted teams of the DARPA project (indicated by abbreviations of their approach).

These results are a significant step forward for DCOP techniques. Our mapping and algorithm successfully solves real-world problems orders of magnitude larger and more complex than previous DCOP applications. Our approach comes very close to the solutions of the DARPA teams, and shows promise for future improvements in this domain and applicability of DCOP to other real-world domains.

6. REFERENCES

- [1] E. Bowring, J. P. Pearce, C. Portway, M. Jain, and M. Tambe. On K-optimal distributed constraint optimization algorithms: new bounds and algorithms. In *AAMAS '08*, 2008.
- [2] A. Farinelli, A. Rogers, A. Petcu, and N. Jennings. Decentralised coordination of low-power embedded devices using the max-sum algorithm. In *AAMAS '08*, 2008.
- [3] B. Horling et. al. The TAEMS White Paper, 1999.
- [4] R. Junges and A. L. C. Bazzan. Evaluating the performance of DCOP algorithms in a real world, dynamic problem. In *AAMAS '08*, 2008.
- [5] A. Petcu and B. Faltings. DPOP: A scalable method for multiagent constraint optimization. In *IJCAI '05*, 2005.