Exploiting the Structure of Distributed Constraint Optimization Problems*

(Doctoral Consortium)

Ferdinando Fioretto
ffiorett@cs.nmsu.edu
Computer Science Department & Mathematics and Computer Science Department
New Mexico State University & University of Udine
Las Cruces, NM 88003, USA

ABSTRACT

In the proposed thesis, we study *Distributed Constraint Optimization Problems (DCOPs)*, which are problems where several agents coordinate with each other to optimize a global cost function. The use of DCOPs has gained momentum, due to their capability of addressing complex and naturally distributed problems. However, the adoption of DCOP on large problems faces two main limitations: *(1) Modeling limitations*, as current resolution methods detach the model from the resolution process, assuming that each agent controls a single variable of the problem; and *(2) Solving capabilities*, as the inability of current approaches to capitalize on the presence of structural information which may allow incoherent/unnecessary data to reticulate among the agents as well as to exploit structure of the agent's local problems. The purpose of the proposed dissertation is to address such limitations, studying how to adapt and integrate insights gained from centralized solving techniques in order to enhance DCOP performance and scalability, enabling their use for the resolution of real-world complex problems. To do so, we hypothesize that one can exploit the DCOP structure in both problem modeling and problem resolution phases.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed AI

Keywords

DCOP; CP; Smart Grid

1. PROGRESS TO DATE

Exploiting Structure from Problem Modeling

Modeling many real-world problems as DCOPs often require each agent to control a large number of variables. However, most DCOP resolution approaches assume that each agent controls exclusively a single variable of the problem. As such, researchers have proposed a number of pre-processing techniques to reformulate DCOPs with multi-variable agents into DCOPs with single-variable agents [8]. Unfortunately, these techniques do not scale with the size of the problem due to their inefficient communication requirements. Therefore, we proposed a DCOP problem decomposition that defines a clear separation between the distributed DCOP resolution and the centralized agent sub-problem resolution. This separation exploits co-locality of agent's variables, allowing the adoption of efficient centralized techniques to solve agent sub-problems, while preserving agent's privacy. Agents coordination is achieved employing a global DCOP algorithm [4]. Using such problem decomposition, allows us to significantly reduce the time of the DCOP resolution process. In addition, the knowledge acquired from the DCOP model allows us to reduce the algorithms communication requirements, when compared to existing pre-processing techniques—which ignore the structural information dictated by the model.

The separation between the DCOP resolution process and the centralized agent problem enabled agents to solve their local problem trough a variety of techniques. Motivated by the high complexity of the agent local problem, we proposed the use of hierarchical parallel models, where each agent can *(i)* solve its local problem independently from those of other agents, and *(ii)* parallelize the computations within its own local problem. Such model builds on top of algorithm-specific characteristics, and may substantially reduces the runtime for several DCOP algorithms classes. For instance, in [2], we suggest to solve independent local problems, in parallel, harnessing the multitude of computational units offered by GPGPUs, which led to significant improvements in the runtime of the algorithm resolution.

These results validate our hypothesis that one can exploit the information encoded in the DCOP model through the use of centralized solutions.

Exploiting Structure from Problem Solving

A number of multi-agent systems require agents to run on battery-powered devices and communicate over wireless networks. This imposes constraints on the number and size of individual messages exchanged among agents. *Inference-based* DCOP algorithms, can be effective in solving such problems. They use techniques from dynamic programming to propagate aggregate information among agents, and while their requirements on the number of messages is linear in the number of agents, their messages have a size that is exponential in the size of the treewidth, which can be up to the number of agents -1. Several works from the DCOP community have recognized the use of hard constraints to reduce the size of

*This research is partially supported by NSF grant HRD-1345232. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the sponsoring organizations, agencies, or the U.S. government.

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

the search space and/or reduce the message size. However, they are limited in exploiting relational information expressed in form of tables and/or associated to the form of domain consistency.

We have contributed to this body of research by introducing a type of consistency, called *Branch Consistency* [3], that applies to paths in pseudo-trees. The effect of enforcing Branch Consistency is the ability to actively exploit hard constraints (either explicitly provided in the problem specification or implicitly described in constraints cost tables) to prune the search space and to reduce the size of the messages exchanged among agents. Such form of consistency enforces a more effective pruning than those based on domain-consistency, leading enhanced efficiency and scalability.

These results validate our hypothesis that centralized reasoning can be adapted to exploit the structure of DCOPs during problem solving.

2. PROPOSED PLAN FOR THE FUTURE

Efficient Local Search Strategies for DCOPs

In the work conducted so far we adapted centralized reasoning techniques to complete DCOP algorithms. Nevertheless, solving DCOPs optimally is NP-hard, therefore for large problems, incomplete DCOP algorithms are desirable. Current incomplete DCOP algorithms have combinations of the following limitations: *(i)* they find local minima without quality guarantees; *(ii)* they provide loose quality assessment (such us those in the class of k-optimality [7]); or *(iii)* they do not exploit problem structures, such as hard constraints.

Therefore, capitalizing on strategies from the centralized constraint reasoning community, we propose to adapt the *Large Neighboring Search strategy (LNS)* [5] to the DCOP resolution process. This technique allows to rapidly find solutions by fixing the variable assignments of a set of agents while optimizing over the others. We believe that LNS is a desirable candidate to DCOP local search because *(a)* emulating the centralized results, it can quickly find local minima, *(b)* it inherently uses insights from the CP techniques to take advantage on the presence of hard constraints, and to refine the solution quality—by constraining the solution bound during the resolution process—and *(c)* it is amenable to parallelization (e.g., if groups of agents can explore several neighbors at a time [1]). We plan in studying the use of machine learning techniques to select which set of variables (agents) to unlock during the DCOP solving phase, as well as ensuring quality guarantees on the solution found. This can be achieved by iteratively solving relaxed versions of the problem (e.g., where only few agents are free to make local moves, and acting on a subset of the DCOP cost functions) and using such solution to retrieve bounds, in the relaxed and in original DCOP.

Distributed Simulator, Modeling Language and Application to the Smart Grid Problems

Despite the wide applicability of the DCOP model, there is no general language being used to formally specify a DCOP. By and large, most stand-alone algorithms specify DCOPs in an ad-hoc manner. Moreover, current DCOP simulators model agents as entities running on the same single machine. Therefore we propose a new agent-based modeling language, which extends the MiniZinc language [6]. Such a DCOP language is (a) more expressive than other adopted formalisms (such as XML-based DCOP descriptions) and (b) it allows the expression of constraints succinctly, in the form of rules, using a well adopted semantics from the constraint reasoning community and allows a fine integration with agents' centralized solvers. Our preliminary results show that such a representation may significantly affect performance, due to the stronger inference that may be derived from explicit constraint representation.

We are also implementing a DCOP solver that uses agent distributed over different machines, and can communicate using several network standard communication protocols. We believe that this is a valuable contribution, as current DCOP simulators suffer from strong communication assumptions (e.g., they assume the same cost for all communications, and direct communications, with no routing), which may not fully reflect the DCOP algorithm behavior on real scenarios.

Lastly, we plan to apply the techniques produced in the proposed thesis on smart grid domains. In particular we are interested in studying agent-based demand-side management problems, which includes a set of *consumers* such as, residential and commercial buildings, and a set of *energy providers*. The providers are in charge of supplying electricity to the consumers, satisfying their load requirements, while minimizing the overall amount of pollutants emitted. The consumers, in turn, may control small generators such as, photovoltaic power generators, storage devices, and electric vehicles, which can be employed to aid in reducing their energy consumption costs. The goal is to help customers make autonomous decisions on their energy consumption and storage profiles, in order to minimize their consumptions costs, while helping the providers reducing peaks in load demands, as well as the overall amount of pollutant emitted. Privacy is one of the core motivation behind the adoption of a DCOP-based model for such problems. However, several challenges arises due to the inability of DCOP to model some form of uncertainty (e.g., derived from the users consumption and generation levels) and environment evolution (e.g., transmission line faults may affect the network topology). Thus, to cope with such limitations, we plan to study hybrid approaches derived by merging solutions from Stochastic Optimization and Decision Theory domains with DCOP ones.

REFERENCES

[1] F. Campeotto, A. Dovier, F. Fioretto, and E. Pontelli. A GPU implementation of large neighborhood search for solving constraint optimization problems. In *ECAI*, pages 189–194, 2014.

[2] F. Fioretto, F. Campeotto, L. D. R. Fioretto, W. Yeoh, and E. Pontelli. GD-Gibbs: A GPU-based sampling algorithm for solving distributed constraint optimization problems (Extended Abstract). In *AAMAS*, pages 1339–1340, 2014.

[3] F. Fioretto, T. Le, W. Yeoh, E. Pontelli, and T. C. Son. Improving DPOP with branch consistency for solving distributed constraint optimization problems. In *CP*, pages 307–323, 2014.

[4] F. Fioretto, W. Yeoh, and E. Pontelli. Decomposition techniques for dcops to exploit multi-variable agents and multi-level parallelism (Extended Abstract). In *AAMAS*, 2015.

[5] P. V. Hentenryck and L. Michel. *Constraint-based local search*. The MIT Press, 2009.

[6] N. Nethercote, P. J. Stuckey, R. Becket, S. Brand, G. J. Duck, and G. Tack. Minizinc: Towards a standard cp modelling language. In *CP*, pages 529–543. 2007.

[7] J. Pearce and M. Tambe. Quality guarantees on k-optimal solutions for distributed constraint optimization problems. In *IJCAI*, pages 1446–1451, 2007.

[8] M. Yokoo, editor. *Distributed Constraint Satisfaction: Foundation of Cooperation in Multi-agent Systems*. Springer, 2001.