Multi-Agent Coordination: DCOPs and Beyond*

Marc Pujol-Gonzalez
Artificial Intelligence Research Institute (IIIA-CSIC)
Universitat Autònoma de Barcelona
Campus UAB, E-08193 Bellaterra, Catalonia (Spain)
mpujol@iiia.csic.es

Abstract

Distributed constraint optimization problems (DCOPs) are a model for representing multi-agent systems in which agents cooperate to optimize a global objective. The DCOP model has two main advantages: it can represent a wide range of problem domains, and it supports the development of generic algorithms to solve them. Firstly, this paper presents some advances in both complete and approximate DCOP algorithms. Secondly, it explains that the DCOP model makes a number of unrealistic assumptions that severely limit its range of application. Finally, it points out hints on how to tackle such limitations.

1 Distributed constraint optimization

Distributed constraint optimization problems (DCOPs) are a model for representing multi-agent systems in which agents cooperate to optimize a global objective. The DCOP model has two main advantages. Firstly, it can represent a wide range of problem domains such as wireless sensor networks [Zhang *et al.*, 2005], peer-to-peer networks [Faltings *et al.*, 2006], meeting scheduling [Maheswaran *et al.*, 2004], and traffic control [Junges and Bazzan, 2008]. Secondly, it supports the development of generic solving algorithms.

Therefore, researchers have developed several complete algorithms such as ADOPT [Modi *et al.*, 2005], DPOP [Petcu and Faltings, 2005], and its generalization GDL [Aji and McEliece, 2000; Vinyals *et al.*, 2010b]. Nevertheless, DCOPs are shown to be NP-Hard [Modi *et al.*, 2005]. Being complete, the main advantage of these algorithms is that they guarantee the maximum possible quality: optimality. However, they scale poorly when the number of agents increases, regarding both computational and communication requirements. Function filtering is a promising technique [Brito and Meseguer, 2010] to achieve better scalability. Basically, given a method to compute approximations of cost functions and a candidate solution, function filtering allows pruning regions of the solution space that only contain non-optimal solutions.

Notice that some application domains are specially communication constrained, whereas others are mainly computationally constrained. For instance, data transmission is severely limited in wireless sensor networks, and bandwidth is a scarce resource in peer-to-peer networks. Conversely, meeting scheduling and traffic control are usually computationally constrained because they mainly operate over high-speed networks. Such distinctions between resource-constrained settings motivated us to study different function approximation methods to be employed along with function filtering, to reduce either communication or computation requirements as much as possible.

First, in [Pujol-Gonzalez *et al.*, 2011] we presented a novel class of approximation techniques, the so-called top-down approximations. Combining these new techniques with function filtering, we managed to reduce communication requirements by as much as two orders of magnitude, while keeping computational requirements at bay. As a consequence, the resulting algorithm appears as a very good candidate to solve DCOPs optimally in communication-constrained scenarios.

Currently, we are working on improving the effectiveness of function filtering in computationally-constrained settings. Since function filtering's pruning is based on lower and upper bounds on the optimal solution cost, tightening such bounds increases the amount of pruning. Because more pruning means a reduction on the solution space to explore, agents require less computational resources (both CPU and memory) to solve the same problem. Likewise, these savings in computational resources also increase the range of problems that can be solved optimally by algorithms employing function filtering. In fact, preliminary results indicate that our improvements allow agents to solve up to 75% more problem instances given the same resource constraints.

Another approach to improve the scalability of DCOP algorithms is to drop optimality in favor of lower complexity, approximate algorithms. Traditionally, these algorithms have not offered any quality guarantees at all [Zhang *et al.*, 2005], but recent works have been able to provide offline bounds for some of them [Farinelli *et al.*, 2009; Kiekintveld *et al.*, 2010]. The disadvantage of such offline bounds is that they are generally very weak. Thus, we provided a new class of al-

*This work has been funded by projects EVE (TIN2009-14702-C02-01 and 02), Agreement Technologies (CONSOLIDER CSD2007-0022), RECEDIT (TIN2009-13591-C02-02) and Generalitat de Catalunya (2009-SGR-1434 and 2009-SGR-362). Marc Pujol-Gonzalez is supported by the Ministry of Science and Innovation (BES-2010-030466)

gorithms that is able to provide (much better) online bounds in [Vinyals *et al.*, 2010a], the so-called Divide and Coordinate (DaC) approach. In DaC, agents divide an intractable DCOP into simpler, tractable, subproblems to individually solve them. Thereafter, agents try solve the DCOP by searching for an agreement on the optimal assignments of their subproblems. The advantages of the DaC approach are two-fold: (1) It has great scalability because it is based on local interactions; and (2) at any time, it provides an online bound on the quality of the current solution.

2 Beyond DCOPs

As mentioned above, the DCOP formalism is useful to model multiple domains. However, it makes a number of unrealistic assumptions that severely limit its range of application.

Firstly, it assumes that agents operate in a static environment. Thus, agents' utilities do not change while the problem is being solved. Moreover, agents' actions are only applied once the problem is solved. Against this background, we plan to work on developing dynamic algorithms that can incorporate both changes in the DCOP structure (new agents joining in and agents' failing) and new information during the solving process. The constant inflow of new information would disrupt the operation of most DCOP algorithms. Despite that, both the GDL with function filtering and the DaC approaches are bound-based algorithms. This is, they operate by computing a lower and an upper bound that progressively get closer and closer till they eventually converge on the optimal solution. Therefore, we plan to extend both approaches so that they accommodate any new information by conveniently loosening their bounds, while still operating correctly and providing good quality guarantees.

Secondly, the DCOP formalism assumes that the environment is deterministic and fully-observable, meaning that agents have complete information about the utility of the outcomes of their possible decisions. Nonetheless, there are many domains of an explorative nature, where one of the agents' objectives is to acquire knowledge about their environment and properly adjust to it. For instance, [Jain *et al.*, 2009] presents a domain where some mobile sensors have to coordinate to establish (and maintain) a wireless network as stable and reliable as possible. Although there are some preliminary works in this direction [Stranders, 2010], they mainly focus on handling the uncertainty of a very specific domain. Hence, we also plan to work on a generalized model of DCOPs that incorporates unreliable and unknown information. Thereafter, we will study whether existing (both complete and approximate) algorithms can be adapted to solve this new class of problems, paying special attention to the exploration-exploitation tradeoff inherent to them.

References

[Aji and McEliece, 2000] Srinivas M. Aji and Robert J. McEliece. The generalized distributive law. *IEEE Transactions on Information Theory*, 46(2):325–343, 2000.

[Brito and Meseguer, 2010] Ismel Brito and Pedro Meseguer. Cluster tree elimination for distributed constraint optimization with quality guarantees. *Fundamenta Informaticae*, 102:263–286, 2010.

[Faltings *et al.*, 2006] B. Faltings, D. Parkes, A. Petcu, and J. Shneidman. Optimizing streaming applications with selfinterested users using M-DPOP. In *COMSOC'06*, pages 206–219, 2006.

[Farinelli *et al.*, 2009] A. Farinelli, A. Rogers, and N. Jennings. Bounded approximate decentralised coordination using the max-sum algorithm. In *IJCAI DCR Workshop*, pages 46–59, 2009.

[Jain *et al.*, 2009] M. Jain, M. Taylor, M. Tambe, and M. Yokoo. DCOPs meet the real world: Exploring unknown reward matrices with applications to mobile sensor networks. In *IJCAI*, pages 181–186, 2009.

[Junges and Bazzan, 2008] Robert Junges and Ana L. C. Bazzan. Evaluating the performance of DCOP algorithms in a real world, dynamic problem. In *AAMAS (2)*, pages 599–606, 2008.

[Kiekintveld *et al.*, 2010] C. Kiekintveld, Z. Yin, A. Kumar, and M. Tambe. Asynchronous algorithms for approximate distributed constraint optimization with quality bounds. In *AAMAS*, pages 133–140, 2010.

[Maheswaran *et al.*, 2004] Rajiv T. Maheswaran, Milind Tambe, Emma Bowring, Jonathan P. Pearce, and Pradeep Varakantham. Taking dcop to the real world: Efficient complete solutions for distributed multi-event scheduling. In *AAMAS*, pages 310–317, 2004.

[Modi *et al.*, 2005] Pragnesh Jay Modi, Wei-Min Shen, Milind Tambe, and Makoto Yokoo. Adopt: asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence*, 161(1-2):149–180, 2005.

[Petcu and Faltings, 2005] Adrian Petcu and Boi Faltings. A scalable method for multiagent constraint optimization. In *IJCAI*, pages 266–271, 2005.

[Pujol-Gonzalez *et al.*, 2011] Marc Pujol-Gonzalez, Jesus Cerquides, Pedro Meseguer, and Juan A. Rodriguez-Aguilar. Communication-constrained DCOPs: Message approximation in GDL with function filtering. In *AAMAS*, 2011.

[Stranders, 2010] R. Stranders. *Decentralised coordination of information gathering agents*. PhD thesis, University of Southampton, 2010.

[Vinyals *et al.*, 2010a] M. Vinyals, M. Pujol, JA Rodriguez-Aguilar, and J. Cerquides. Divide-and-Coordinate: DCOPs by agreement. In *AAMAS*, pages 149–156, 2010.

[Vinyals *et al.*, 2010b] Meritxell Vinyals, Juan A. Rodríguez-Aguilar, and Jesús Cerquides. Constructing a unifying theory of dynamic programming dcop algorithms via the generalized distributive law. *JAAMAS*, pages 1–26, 2010.

[Zhang *et al.*, 2005] Weixiong Zhang, Guandong Wang, Zhao Xing, and Lars Wittenburg. Distributed stochastic search and distributed breakout: properties, comparison and applications to constraint optimization problems in sensor networks. *Artif. Intell.*, 161(1-2):55–87, 2005.