AI Buzzwords Explained: Distributed Constraint Optimization Problems

Ferdinando Fioretto (University of Michigan; fioretto@umich.edu)
William Yeoh (Washington University in St. Louis; wyeoh@wustl.edu)
DOI: 10.1145/3175502.3175506

The power network is the largest operating *machine* on earth, generating more than US$400bn a year[1] keeping the lights on for our homes, offices, and factories. A significant concern in power networks is for the energy providers to be able to generate enough power to supply the demands at any point in time. Short terms demand peaks are however hard to predict and, thus, in the modern *smart electricity grid*, the energy providers can exploit the demand-side flexibility of the consumers to reduce the peaks in load demand.

This control mechanism is called *Demand-side management* (DSM). DSM can be obtained by scheduling *shiftable loads* (i.e., a portion of power consumption that can be moved from a time slot to another) from peak to off-peak hours (Fioretto, Yeoh, & Pontelli, 2017; Logenthiran, Srinivasan, & Shun, 2012; Voice, Vytelingum, Ramchurn, Rogers, & Jennings, 2011). In a simplified version of this problem, the energy provider has a desired maximal amount of power that it can generate and, thus, use to serve its customers. When the predicted amount of customer loads exceed such amount, the provider has to reschedule some of these loads in different time slots to satisfy the constraint on the maximum power capacity.

Such an approach, however, requires the provider to control a portion of the consumer's electrical appliances, affecting privacy and users' autonomy. On the other hand, residential and commercial buildings are progressively being partially automated, through the introduction of smart devices (e.g., smart thermostats, circulator heating, washing machines). Household penetration is at 5.8% in 2016 and is expected to hit 18.6% in 2020.[2]

Device scheduling can be executed by the users, without the control of a centralized authority, and a coordinated device scheduling within a neighborhood of buildings can be used as a DSM strategy, preserving user data privacy. Figure 1 illustrates such scenario.

One possible way to solve this problem is through the use of *Distributed Constraint Optimization Problems* (DCOPs) (Fioretto, Pontelli, & Yeoh, 2016; Modi, Shen, Tambe, & Yokoo, 2005; Petcu & Faltings, 2005a). DCOP algorithms are a class of distributed cooperative multi-agent algorithms in which several autonomous agents coordinate their decisions to achieve a shared goal while accounting for personal preferences. The agents can be thought of as software programs whose execution does not depend on the execution of other agents. Their actions are expressed using the concept of *variables*, i.e., abstract entities that can take one out of several values (describing the possible set of actions for the agent). Each agent needs to decide which value to assign to its variables. The outcome of an action is expressed in terms of a *cost* (or *reward*) and typically depends on the joint action of multiple agents. The goal of a DCOP is expressed in the form of an objective function to be minimized (or maximized). To coordinate their actions, agents employ a message passing mechanism realized through a networked communication.

Mathematically, a DCOP is composed by the following entities:

- $\mathcal{A} = \{a_1, \dots, a_p\}$: The set of autonomous agents participating in the problem.
- $\mathcal{X} = \{x_1, \dots, x_n\}$: The set of variables in the problem. Each variable is controlled by exactly one agent.
- $\mathcal{D} = \{D_1, \dots, D_n\}$: The domains for the variables in $\mathcal{X}$, where D_i represents the set of possible values that the variable x_i may be assigned.

[1] U.S. Energy Information Administration

[2] https://www.statista.com/outlook/279/109/smart-home/united-states#market-driver

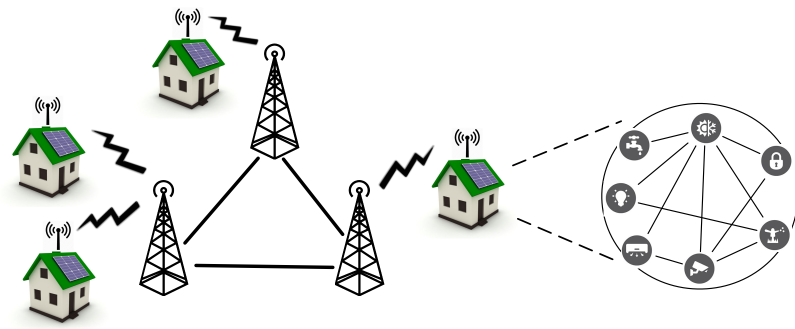

Figure 1: An illustration of a smart neighborhood with each home controlling a set of smart devices.

- $\mathcal{C} = \{c_1, \dots, c_e\}$: The set of problem constraints. Each constraint c_i is a function that involves (multiple) variable(s) from $\mathcal{X}$ and associates a cost for each combination of their value assignments.
- $\alpha : \mathcal{X} \rightarrow \mathcal{A}$: A mapping that associates variables to agents, expressing which agent controls which variables.

The goal of the problem is to find an assignment for the agent variables that minimizes the sum of all costs over all constraints. Since the agents are physically distributed across a network, all communication take the form of messages. Thus, agents coordinate the value assignment for their variables following a given distributed protocol. In addition, agents knowledge is limited to their resources: each agent knows exclusively the outcomes of the variables it controls and the constraints it shares with some other agents. This scheme is effective to design algorithms that preserve data privacy.

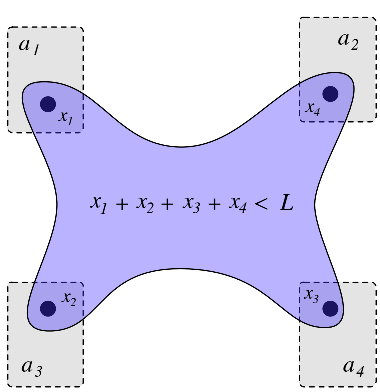

Figure 2: An illustration of a smart neighbor with each home controlling a set of smart devices.

In the DSM smart device scheduling example in Figure 2, we illustrate a smart neighborhood composed of 4 homes, each represented by one agent: $a_1, \dots, a_4$. In this simplified scenario, each agent controls one single variable $x_1, \dots, x_4$, which represents a simplified electrical appliance that can be switched *off* (consuming 0 kWh) or *on* (consuming 2 kWh). Assume the energy provider imposes a total consumption limit of 4 kWh. We represent the domains of each variable with the set $\{0, 2\}$, indicating the consumption associated with the device's actions. The only constraint of the problem is expressed with the formula $x_1 + x_2 + x_3 + x_4 \leq 4$, meaning that the aggregated energy consumption cannot exceed 4 kWh. A possible solution to the problem is thus having agents a_1, and a_2 switching their appliances *on*, and thus consuming a total of $x_1 = 2 + x_2 = 2 = 4$ kWh, and agents a_3 and a_4 switching their appliance *off*, thus consuming 0 kWh. A more detailed description of the DSM scheduling application and its corresponding DCOP model can be found in (Fioretto et al., 2017; Tabakhi, Le, Fioretto, & Yeoh, 2017).

The DCOP framework is general and offers a flexible tool to model a wide variety of problems. Examples of use of DCOPs to solve distributed problems include *service-oriented computing*, that relies on sharing resources over a network, focusing on maximizing the effectiveness of the shared resources which are used by multiple applications (Choudhury, Dey, Dutta, & Choudhury, 2014; Jin, Cao, & Li, 2011; Li, Wang, Ding, & Li, 2014), *sensor network problems*, which consist of coordinating a large number of inexpensive and autonomous sensor nodes, constrained by a limited communication range and battery life (Hosseini Semnani & Basir, 2013; Ota, Matsui, & Matsuo, 2009; Stranders, Farinelli, Rogers, & Jennings, 2009; Zhang, Wang, Xing, & Wittenberg, 2005), and many others (Brys, Pham, & Taylor, 2014; Gaudreault, Frayret, & Pesant, 2009; Junges & Bazzan, 2008; Kumar, Faltings, & Petcu, 2009; Miller, Ramchurn, & Rogers, 2012; Rust, Picard, & Ramparany, 2016; Yeoh & Yokoo, 2012; Zivan, Yedidsion, Okamoto, Glinton, & Sycara, 2015). More examples can be found in a recent survey (Fioretto, Pontelli, & Yeoh, 2016).

It turns out that it is difficult (NP-hard) to optimally solve this kind of problems, and that there is a close relationship between the amount of information that needs to be encoded in a message (message size) vs. the number of messages exchanged by the agents (network load). Thus, an extensive piece of the DCOP literature focuses on the study of algorithms that trade off solution quality for faster runtime and reduced use of network resources (Farinelli, Rogers, Petcu, & Jennings, 2008; Fioretto, Yeoh, & Pontelli, 2016; Maheswaran, Pearce, & Tambe, 2004; Nguyen, Yeoh, & Lau, 2013; Ottens, Dimitrakakis, & Faltings, 2017; Pearce & Tambe, 2007; Petcu & Faltings, 2007a; Yeoh, Sun, & Koenig, 2009; Zhang et al., 2005).

As one of the motivations for the use of DCOPs is the preservation of privacy, there is also a large body of work on privacy-preserving algorithms (Grinshpoun & Tassa, 2016; Léauté & Faltings, 2011a, 2013; Tassa, Grinshpoun, & Zivan, 2017; Tassa, Zivan, & Grinshpoun, 2016).

Additionally, the DCOP model has also been extended to handle problems where agents have multiple objectives (Delle Fave, Stranders, Rogers, & Jennings, 2011; Matsui, Silaghi, Hirayama, Yokoo, & Matsuo, 2012), problems of a dynamic nature (i.e., where the problem changes over time) (Hoang et al., 2016, 2017; Nguyen, Yeoh, Lau, Zilberstein, & Zhang, 2014; Petcu & Faltings, 2005b, 2007b; Yeoh, Varakantham, Sun, & Koenig, 2015), and problems with uncertainty (i.e., where constraint costs depend from uncertain factors, such as weather) (Atlas & Decker, 2010; Le, Fioretto, Yeoh, Son, & Pontelli, 2016; Léauté & Faltings, 2011b; Nguyen, Yeoh, & Lau, 2012; Stranders, Delle Fave, Rogers, & Jennings, 2011).

A forum for discussion on DCOP algorithms and applications has been held at the *Optimization in Multi-Agent Systems* (OptMAS) workshop since 2010 at the *International Conference on Autonomous Agents and Multiagent Systems* (AAMAS). Recent dissertations within the past 5 years include (Billiau, 2015; Delle Fave, 2012; Fioretto, 2016; Grubshtein, 2012; Guttierez, 2012; Hanada, 2017; Hatano, 2013; Kim, 2015; Miller, 2014; Netzer, 2015; Okimoto, 2012; Ottens, 2012; Ueda, 2014; Yedidsion, 2015).

References

Atlas, J., & Decker, K. (2010). Coordination for uncertain outcomes using distributed neighbor exchange. In *Proceedings of the international conference on autonomous agents and multiagent systems* (pp. 1047–1054).

Billiau, G. (2015). *An equitable approach of solving distributed constraint optimization problems* (PhD thesis). University of Wollongong, Wollongong (Australia).

Brys, T., Pham, T. T., & Taylor, M. E. (2014). Distributed learning and multi-objectivity in traffic light control. *Connection Science*, *26*(1), 65–83.

Choudhury, B., Dey, P., Dutta, A., & Choudhury, S. (2014). A multi-agent based optimised server selection scheme for SOC in pervasive environment. In *Advances in practical applications of heterogeneous multi-agent systems* (pp. 50–61).

Delle Fave, F. M. (2012). *Theory and practice of coordination algorithms exploiting the generalised distributive law* (PhD thesis). University of Southampton, Southampton (UK).

Delle Fave, F. M., Stranders, R., Rogers, A., & Jennings, N. (2011). Bounded decentralised coordination over multiple objectives. In *Proceedings of the international conference on autonomous agents and multiagent systems* (pp. 371–378).

Farinelli, A., Rogers, A., Petcu, A., & Jennings, N. (2008). Decentralised coordination of low-power embedded devices using the Max-Sum algorithm. In *Proceedings of the international conference on autonomous agents and multiagent systems* (pp. 639–646).

Fioretto, F. (2016). *Exploiting the structure of distributed constraint optimization problems with applications in smart grids* (PhD thesis). New Mexico State University, Las Cruces (United States).

Fioretto, F., Pontelli, E., & Yeoh, W. (2016). Distributed constraint optimization problems and applications: A survey. *CoRR*, *abs/1602.06347*.

Fioretto, F., Yeoh, W., & Pontelli, E. (2016). Multi-variable agent decomposition for DCOPs. In *Proceedings of the AAAI conference on artificial intelligence* (pp. 2480–2486).

Fioretto, F., Yeoh, W., & Pontelli, E. (2017). A multiagent system approach to scheduling devices in smart homes. In *Proceedings of the international conference on autonomous agents and multiagent systems* (pp. 981–989).

Gaudreault, J., Frayret, J.-M., & Pesant, G. (2009). Distributed search for supply chain coordination. *Computers in Industry*, *60*(6), 441–451.

Grinshpoun, T., & Tassa, T. (2016). P-SyncBB: A privacy preserving branch and bound DCOP algorithm. *Journal of Artificial Intelligence Research*, *57*, 621–660.

Grubshtein, A. (2012). *Distribute search by agents with personal preferences* (PhD thesis). Ben-Gurion University of the Negev, Beer-Sheva (Israel).

Guttierez, P. (2012). *Distributed constraint optimization related with soft arc consistency* (PhD thesis). Universitat Autònoma de Barcelona, Bellaterra (Spain).

Hanada, K. (2017). *Distributed lagrangian relaxation protocol and its applications* (PhD thesis). Kobe University, Kobe (Japan).

Hatano, D. (2013). *Solutions for constraint problem under the dynamic/distributed environment* (PhD thesis). Kobe University, Kobe (Japan).

Hoang, K. D., Fioretto, F., Hou, P., Yokoo, M., Yeoh, W., & Zivan, R. (2016). Proactive dynamic distributed constraint optimization. In *Proceedings of the international conference on autonomous agents and multiagent systems* (pp. 597–605).

Hoang, K. D., Hou, P., Fioretto, F., Yeoh, W., Zivan, R., & Yokoo, M. (2017). Infinite-horizon proactive dynamic DCOPs. In *Proceedings of the international conference on autonomous agents and multiagent systems* (pp. 212–220).

Hosseini Semnani, S., & Basir, O. A. (2013). Target to sensor allocation: A hierarchical dynamic distributed constraint optimization approach. *Computer Communications*, *36*(9), 1024–1038.

Jin, Z., Cao, J., & Li, M. (2011). A distributed application component placement approach for cloud computing environment. In *Proceedings of the international conference on dependable, autonomic and secure computing* (pp. 488–495).

Junges, R., & Bazzan, A. L. (2008). Evaluating the performance of DCOP algorithms in a real world, dynamic problem. In *Proceedings of the international conference on autonomous agents and multiagent systems* (pp. 599–606).

Kim, Y. (2015). *Application of techniques for MAP estimation to distributed constraint optimization problem* (PhD thesis). University of Massachusetts at Amherst, Amherst (United States).

Kumar, A., Faltings, B., & Petcu, A. (2009). Distributed constraint optimization with structured resource constraints. In *Proceedings of the international conference on autonomous agents and multiagent systems* (pp. 923–930).

Le, T., Fioretto, F., Yeoh, W., Son, T. C., & Pontelli, E. (2016). 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* (pp. 606–614).

Léauté, T., & Faltings, B. (2011a). Coordinating logistics operations with privacy guarantees. In *Proceedings of the international joint conference on artificial intelligence (ijcai)* (pp. 2482–2487).

Léauté, T., & Faltings, B. (2011b). Distributed constraint optimization under stochastic uncertainty. In *Proceedings of the aaai conference on artificial intelligence (aaai)* (pp. 68–73).

Léauté, T., & Faltings, B. (2013). Protecting privacy through distributed computation in multi-agent decision making. *Journal of Artificial Intelligence Research*, *47*, 649–695.

Li, X., Wang, H., Ding, B., & Li, X. (2014). MABP: an optimal resource allocation approach in data center networks. *Science China Information Sciences*, *57*(10), 1–16.

Logenthiran, T., Srinivasan, D., & Shun, T. (2012). Demand side management in smart grid using heuristic optimization. *IEEE Transactions on Smart Grid*, *3*(3), 1244–1252.

Maheswaran, R., Pearce, J., & Tambe, M. (2004). Distributed algorithms for DCOP: A graphical game-based approach. In *Proceedings of the international conference on parallel and distributed computing systems* (pp. 432–439).

Matsui, T., Silaghi, M., Hirayama, K., Yokoo, M., & Matsuo, H. (2012). Distributed search method with bounded cost vectors on multiple objective DCOPs. In *Proceedings of the principles and practice of multi-agent systems* (pp. 137–152).

Miller, S. (2014). *Decentralised coordination of smart distribution networks using message passing* (PhD thesis). University of Southampton, Southampton (UK).

Miller, S., Ramchurn, S. D., & Rogers, A. (2012). Optimal Decentralised Dispatch of Embedded Generation in the Smart Grid. In *Proceedings of the international conference on autonomous agents and multiagent systems* (pp. 281–288).

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

Netzer, A. (2015). *Distributed constraint optimization, from utilitarianism to fairness* (PhD thesis). Ben-Gurion University of the Negev, Beer-Sheva (Israel).

Nguyen, D. T., Yeoh, W., & Lau, H. C. (2012). Stochastic Dominance in Stochastic DCOPs for Risk-sensitive Applications. In *Proceedings of the international conference on autonomous agents and multiagent systems* (pp. 257–264).

Nguyen, D. T., Yeoh, W., & Lau, H. C. (2013). Distributed Gibbs: A memory-bounded sampling-based DCOP algorithm. In *Proceedings of the international conference on autonomous agents and multiagent systems* (pp. 167–174).

Nguyen, D. T., Yeoh, W., Lau, H. C., Zilberstein, S., & Zhang, C. (2014). Decentralized multi-agent reinforcement learning in average-reward dynamic DCOPs. In *Proceedings of the AAAI conference on artificial intelligence* (pp. 1447–1455).

Okimoto, T. (2012). *Graph-based algorithms for distributed constraint satisfaction/optimization problems* (PhD thesis). Kyushu University, Fukuoka (Japan).

Ota, K., Matsui, T., & Matsuo, H. (2009). Layered distributed constraint optimization problem for resource allocation problem in distributed sensor networks. In *Proceedings of the principles and practice of multi-agent systems* (pp. 245–260).

Ottens, B. (2012). *Coordination and sampling in distributed constraint optimization* (PhD thesis). Ecole Polytechnique Fédérale de Lausanne, Lausanne (Switzerland).

Ottens, B., Dimitrakakis, C., & Faltings, B. (2017). DUCT: An upper confidence bound approach to distributed constraint optimization problems. *ACM Transactions on Intelligent Systems and Technology*, *8*(5), 69:1–69:27.

Pearce, J., & Tambe, M. (2007). Quality guarantees on k-optimal solutions for distributed constraint optimization problems. In *Proceedings of the international joint conference on artificial intelligence* (pp. 1446–1451).

Petcu, A., & Faltings, B. (2005a). A scalable method for multiagent constraint optimization. In *Proceedings of the international joint conference on artificial intelligence* (pp. 1413–1420).

Petcu, A., & Faltings, B. (2005b). Superstabilizing, Fault-Containing Distributed Combinatorial Optimization. In *Proceedings of the AAAI conference on artificial intelligence* (pp. 449–454).

Petcu, A., & Faltings, B. (2007a). MB-DPOP: A new memory-bounded algorithm for distributed optimization. In *Proceedings of the international joint conference on artificial intelligence* (pp. 1452–1457).

Petcu, A., & Faltings, B. (2007b). Optimal solution stability in dynamic, distributed constraint optimization. In *Proceedings of the international conference on intelligent agent technology* (pp. 321–327).

Rust, P., Picard, G., & Ramparany, F. (2016). Using message-passing DCOP algo-

rithms to solve energy-efficient smart environment configuration problems. In *Proceedings of the international joint conference on artificial intelligence* (pp. 468–474).
Stranders, R., Delle Fave, F. M., Rogers, A., & Jennings, N. (2011). *U-GDL: A decentralised algorithm for dcops with uncertainty* (Tech. Rep.). Department of Electronics and Computer Science: University of Southampton.
Stranders, R., Farinelli, A., Rogers, A., & Jennings, N. R. (2009). Decentralised coordination of continuously valued control parameters using the max-sum algorithm. In *Proceedings of the international conference on autonomous agents and multiagent systems* (pp. 601–608).
Tabakhi, A. M., Le, T., Fioretto, F., & Yeoh, W. (2017). Preference elicitation for DCOPs. In *Proceedings of the international conference on principles and practice of constraint programming* (pp. 278–296).
Tassa, T., Grinshpoun, T., & Zivan, R. (2017). Privacy preserving implementation of the Max-Sum algorithm and its variants. *Journal of Artificial Intelligence Research*, *59*, 311–349.
Tassa, T., Zivan, R., & Grinshpoun, T. (2016). Preserving privacy in region optimal DCOP algorithms. In *Proceedings of the international joint conference on artificial intelligence* (pp. 496–502).
Ueda, S. (2014). *Computational coalition formation: Compact representation and constrained matching* (PhD thesis). Kyushu University, Fukuoka (Japan).
Voice, T., Vytelingum, P., Ramchurn, S., Rogers, A., & Jennings, N. (2011). Decentralised control of micro-storage in the smart grid. In *Proceedings of the AAAI conference on artificial intelligence* (pp. 1421–1427).
Yedidsion, H. (2015). *Distributed constraint optimization for teams of mobile agents* (PhD thesis). Ben-Gurion University of the Negev, Beer-Sheva (Israel).
Yeoh, W., Sun, X., & Koenig, S. (2009). Trading off solution quality for faster computation in DCOP search algorithms. In *Proceedings of the international joint conference on artificial intelligence* (pp. 354–360).
Yeoh, W., Varakantham, P., Sun, X., & Koenig, S. (2015). Incremental DCOP search algorithms for solving dynamic DCOP problems. In *Proceedings of the international conference on intelligent agent technology* (pp. 257–263).
Yeoh, W., & Yokoo, M. (2012). Distributed problem solving. *AI Magazine*, *33*(3), 53–65.
Zhang, W., Wang, G., Xing, Z., & Wittenberg, L. (2005). Distributed stochastic search and distributed breakout: Properties, comparison and applications to constraint optimization problems in sensor networks. *Artificial Intelligence*, *161*(1–2), 55–87.
Zivan, R., Yedidsion, H., Okamoto, S., Glinton, R., & Sycara, K. (2015). Distributed constraint optimization for teams of mobile sensing agents. *Journal of Autonomous Agents and Multi-Agent Systems*, *29*(3), 495–536.

Ferdinando Fioretto is a postdoctoral researcher at the Industrial and Operations Engineering Department of the University of Michigan. His research focuses on multi-agent systems, data privacy, and discrete optimization. His dissertation was awarded the best dissertation in Artificial Intelligence from the Italian Association of Artificial Intelligence, in 2017. Additional information can be found at: http://www-personal.umich.edu/~fioretto/.

William Yeoh is an assistant professor in the Computer Science and Engineering Department at Washington University in St. Louis. His research interests include multi-agent systems, distributed constraint reasoning, heuristic search, and planning with uncertainty. He is an NSF CAREER awardee and was named in IEEE's AI's 10-to-Watch list in 2015. Additional information can be found at: https://sites.wustl.edu/wyeoh/.