Distributed Constraint Optimization Problems and Applications: A Survey

Ferdinando Fioretto   FIORETTO@UMICH.EDU
Department of Industrial and Operations Engineering
University of Michigan
Ann Arbor, MI 48109, USA

Enrico Pontelli   EPONTELL@CS.NMSU.EDU
Department of Computer Science
New Mexico State University
Las Cruces, NM 88003, USA

William Yeoh   WYEOH@WUSTL.EDU
Department of Computer Science and Engineering
Washington University in St. Louis
St. Louis, MO 63130, USA

Abstract

The field of *multi-agent system (MAS)* is an active area of research within artificial intelligence, with an increasingly important impact in industrial and other real-world applications. In a MAS, autonomous agents interact to pursue personal interests and/or to achieve common objectives. *Distributed Constraint Optimization Problems (DCOPs)* have emerged as a prominent agent model to govern the agents' autonomous behavior, where both algorithms and communication models are driven by the structure of the specific problem. During the last decade, several extensions to the DCOP model have been proposed to enable support of MAS in complex, real-time, and uncertain environments.

This survey provides an overview of the DCOP model, offering a classification of its multiple extensions and addressing both resolution methods and applications that find a natural mapping within each class of DCOPs. The proposed classification suggests several future perspectives for DCOP extensions and identifies challenges in the design of efficient resolution algorithms, possibly through the adaptation of strategies from different areas.

1. Introduction

An *agent* can be defined as an entity, or computer program, that behaves autonomously within an arbitrary system in the pursuit of some goals (Wooldridge, 2009). A *multi-agent system (MAS)* is a system where multiple agents interact in the pursuit of such goals. Within a MAS, agents may interact with each other *directly,* via communication acts, or *indirectly,* by acting on the shared environment. In addition, agents may decide to *cooperate,* to achieve a common goal, or to *compete,* to serve their own interests at the expense of other agents. In particular, agents may form cooperative *teams*, which can in turn compete against other teams of agents. Multi-agent systems play an important role in *distributed artificial intelligence,* thanks to their ability to model a wide variety of real-world scenarios, where information and control are decentralized and distributed among a set of agents.
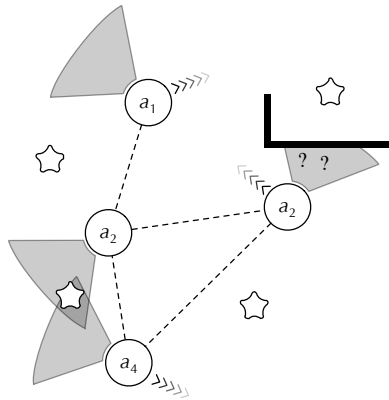
Figure 1: A multi-agent system: Sensors (agents) seek to determine the position of the targets.

Figure 1 illustrates a MAS example. It represents a sensor network where a group of agents, equipped with sensors, seeks to determine the position of some targets. Agents may interact with each other and move away from the current position. The figure depicts the targets as star-shaped objects. The dotted lines define an interaction graph and the directional arrows illustrate agents' movements. In addition, various events that obstruct the sensors of an agent may dynamically occur. For instance, the presence of an obstacle along the agent's sensing range may be detected after the agent's movement.

Within a MAS, an agent is:

- *Autonomous*, as it operates without the direct intervention of humans or other entities and has full control over its own actions and internal state (e.g., in the example, an agent can decide to sense, to move, etc.);

- *Interactant*, in the sense that it interacts with other agents in order to achieve its objectives (e.g., in the example, agents may exchange information concerning results of sensing activities);

- *Reactive*, as it responds to changes that occur in the environment and/or to the requests from other agents (e.g., in the example, agents may react with a move action to the sudden appearance of obstacles).

- *Proactive*, because of its goal-driven behavior, which allows the agent to take initiatives beyond the reactions in response to its environment.

Agent architectures are the fundamental mechanisms underlying the autonomous agent components, supporting their behavior in real-world, dynamic, and uncertain environments. Agent architectures based on *decision theory*, *game theory*, and *constraint programming* have successfully been developed and are popular in the *Autonomous Agents and Multi-Agent Systems (AAMAS)* community.

Decision theory (Raiffa, 1968) assumes that the agent's actions and the environment are inherently uncertain and models such uncertainty explicitly. Agents acting in complex and dynamic environments are required to deal with various sources of uncertainty. The *Decentralized Partially Observable Markov Decision Processes (Dec-POMDPs)* framework (Bernstein, Givan, Immerman, & Zilberstein, 2002) is one of the most general multi-agent frameworks, focused on team coordination in presence of uncertainty about agents' actions and observations. The ability to capture a wide range of complex scenarios makes Dec-POMDPs of central interest within MAS research.

However, the result of this generality is a high complexity for generating optimal solutions. Dec-POMDPs are non-deterministic exponential (NEXP) complete (Bernstein et al., 2002), even for two-agent problems, and scalability remains a critical challenge (Amato, Chowdhary, Geramifard, Ure, & Kochenderfer, 2013).

Game theory (Binmore, 1992) studies interactions between self-interested agents, aiming at maximizing the welfare of the participants. Some of the most compelling applications of game theory to MAS have been in the area of *auctions* and *negotiations* (Kraus, 1997; Noriega & Sierra, 1999; Parsons & Wooldridge, 2002). These approaches model the trading process by which agents can reach agreements on matters of common interest, using market oriented and cooperative mechanisms, such as reaching *Nash equilibria.* Typical resolution approaches aim at deriving a set of equilibrium strategies for each agent, such that, when these strategies are employed, no agent can profit by unilaterally deviating from their strategies. A limitation of game theoretical-based approaches is the lack of an agent's ability to reason upon a global objective, as the underlying model relies on the interactions of self-interested agents.

Constraint programming (Rossi, Beek, & Walsh, 2006) aims at solving decision-making problems formulated as optimization problems of some real-world objective. Constraint programs use the notion of constraints – i.e., relations among entities of the problems (*variables*) – in both problem modeling and problem solving. Constraint programming relies on inference techniques that prevent the exploration of those parts of the solution search space whose assignments to variables are inconsistent with the constraints and/or dominated with respect to the objective function. *Distributed Constraint Optimization Problems (DCOPs)* (Modi, Shen, Tambe, & Yokoo, 2005; Petcu & Faltings, 2005b; Gershman, Meisels, & Zivan, 2009; Yeoh & Yokoo, 2012) are problems where agents need to coordinate their value assignments, in a decentralized manner, to optimize their objective functions. DCOPs focus on attaining a global optimum given the interaction graph of a collection of agents. This approach can be effectively used to model a wide range of problems. Problem solving and communication strategies are directly linked in DCOPs. This feature makes the algorithmic components of a DCOP suitable for exploiting the structure of the interaction graph of the agents to generate efficient solutions.

The absence of a framework to model dynamic problems and uncertainty makes DCOPs unsuitable at solving certain classes of multi-agent problems, such as those characterized by action uncertainty and dynamic environments. However, since its original introduction, the DCOP model has undergone a process of continuous evolution to capture diverse characteristics of agent behavior and the environment in which they operate. Researchers have proposed a number of DCOP frameworks that differ from each other in terms of expressiveness and classes of problem they can target, extending the DCOP model to handle both dynamic and uncertain environments. However, current research has not explored how the different DCOP frameworks relate to each other within the general MAS context, which is critical to understand: *(i)* What resolution methods could be borrowed from other MAS paradigms, and *(ii)* What applications can be most effectively modeled within each framework. While there are important existing surveys for Distributed Constraint Satisfaction (Yokoo & Hirayama, 2000) and Distributed Constraint Optimization (Meisels, 2008), this survey aims to comprehensively analyze and categorize the different DCOP frameworks proposed by the MAS community. We do so by presenting an extensive review of the DCOP model and its extensions, the different resolution methods, as well as a number of applications modeled within each particular DCOP extension. This analysis also provides opportunities to identify open challenges and discuss future directions in the general DCOP research area.

List of key symbols			
a_i	Agent	$\pi(\cdot)$	Projection operator
x_i	Decision variable	$p_i(\cdot)$	Probability function
r_i	Random variable	L_{a_i}	a_i's local variables
D_i	Domain of x_i	N_{a_i}	a_i's neighbors
Ω_i	Event space of r_i	C_{a_i}	a_i's children
f_i	Cost function	PC_i	a_i's pseudo-children
$\mathbf{x}^i$	Scope of f_i	P_{a_i}	a_i's parent
m	Number of agents	PP_{a_i}	a_i's pseudo-parents
n	Number of variables	$\alpha(f_i)$	Agents whose variables are in $\mathbf{x}^i$
q	Number of random variables	E_C	Set of edges of the constraint graph
k	Number of cost functions	E_T	Tree edges of the pseudo-tree
d	Size of the largest domain	E_F	Set of edges of the factor graph
$\mathbf{F}_g$	Global objective function	w^*	Induced width of the pseudo-tree
$\vec{\mathbf{F}}$	Vector of objective functions	l	Size of the largest neighborhood
F_i	Objective function in $\vec{\mathbf{F}}$	z	Size of the largest local variable set
$\vec{\mathbf{F}}^\circ$	Utopia point	s	Maximal sample size
$\bot$	Infeasible value	p	Size of the Pareto set
σ	Complete assignment	b	Size of the largest bin
σ_V	Partial assignment for the variables in $V \subseteq \mathbf{X}$	ℓ	Number of iterations of the algorithm
Σ	State space		

Table 1: Commonly Used Symbols and Notations

This survey paper is organized as follows. The next section provides an overview on two relevant constraint satisfaction models and their generalization to the distributed cases. Section 3 introduces DCOPs, overviews the representation and coordination models adopted during the resolution of DCOPs, and it proposes a classification of the different variants of DCOPs based on the characteristics of the agents and the environment. Section 4 presents the classical DCOP model as well as two notable extensions: One characterized by asymmetric cost functions and another by multi-objective optimization. Section 5 presents a DCOP model where the environment changes over time. Section 6 discusses DCOP models in which agents act under uncertainty and may have partial knowledge of the environment in which they act. Section 7 discusses DCOP models in which agents are non-cooperative. For each of these models, the paper introduces their formal definitions, discusses related concepts, and describes several resolution algorithms. A summary of the various classes of problems discussed in this survey is given in Table 6. Section 8 describes a number of applications that have been proposed in the DCOP literature. Section 9 provides a critical review on the DCOP variants surveyed and focuses on their applicability in various settings. Additionally, it describes some potential future directions for research. Finally, Section 10 provides concluding remarks. To facilitate the reading of this survey, Table 1 summarizes the most commonly used symbols and notations.

2. Overview of (Distributed) Constraint Satisfaction and Optimization

This section provides an overview of several constraint satisfaction models, which form the foundation of DCOPs. Figure 2 illustrates the relations among these models.

2.1 Constraint Satisfaction Problems

Constraint Satisfaction Problems (CSPs) (Golomb & Baumert, 1965; Mackworth & Freuder, 1985; Apt, 2003; Rossi et al., 2006) are decision problems that involve the assignment of values to variables, under a set of specified constraints on how variable values should be related to each other. A number of problems can be formulated as CSPs, including resource allocation, vehicle routing, circuit diagnosis, scheduling, and bioinformatics. Over the years, CSPs have become the paradigm of choice to address difficult combinatorial problems, drawing and integrating insights from diverse domains, including artificial intelligence and operations research (Rossi et al., 2006).

A CSP is a tuple $\langle \mathbf{X}, \mathbf{D}, \mathbf{C} \rangle$, where:

- $\mathbf{X} = \{x_1, \ldots, x_n\}$ is a finite set of variables.
- $\mathbf{D} = \{D_1, \ldots, D_n\}$ is a set of finite domains for the variables in $\mathbf{X}$, with D_i being the set of possible values for the variable x_i.
- $\mathbf{C}$ is a finite set of constraints over subsets of $\mathbf{X}$, where a constraint c_i, defined on the k variables $x_{i_1}, \ldots, x_{i_k}$, is a relation $c_i \subseteq \times_{j=1}^{k} D_{i_j}$, where $\{i_1, \ldots, i_k\} \subseteq \{1, \ldots, n\}$. The set of variables $\mathbf{x}^i = \{x_{i_1}, \ldots, x_{i_k}\}$ is referred to as the *scope* of c_i.[1] c_i is called a *unary* constraint if $k = 1$ and a *binary* constraint if $k = 2$. For all other values of k, the constraint is called a *k*-ary constraint.[2]

A *partial assignment* is a value assignment for a proper subset of variables from $\mathbf{X}$ that is consistent with their respective domains, i.e., it is a partial function $\sigma : \mathbf{X} \rightarrow \bigcup_{i=1}^{n} D_i$ such that, for each $x_j \in \mathbf{X}$, if $\sigma(x_j)$ is defined, then $\sigma(x_j) \in D_j$. An assignment is *complete* if it assigns a value to each variable in $\mathbf{X}$. The notation σ is used to denote a complete assignment, and, for a set of variables $\mathbf{V} = \{x_{i_1}, \ldots, x_{i_h}\} \subseteq \mathbf{X}$, $\sigma_{\mathbf{V}} = \langle \sigma(x_{i_1}), \ldots, \sigma(x_{i_h}) \rangle$ to denote the projection of the values in σ associated to the variables in $\mathbf{V}$, where $i_1 < \cdots < i_h$. The goal in a CSP is to find a complete assignment σ such that, for each $c_i \in \mathbf{C}$, $\sigma_{\mathbf{x}^i} \in c_i$, that is, a complete assignment that satisfies all the problem constraints. Such a complete assignment is called a *solution* of the CSP.

2.2 Weighted Constraint Satisfaction Problems

A solution of a CSP must satisfy all of its constraints. In many practical cases, however, it is desirable to consider complete assignments whose constraints can be violated according to a violation degree. The *Weighted Constraint Satisfaction Problem (WCSP)* (Shapiro & Haralick, 1981; Larrosa, 2002) was introduced to capture this property. WCSPs are problems whose constraints are considered as preferences that specify the extent of satisfaction (or violation) of the associated constraint.

A WCSP is a tuple $\langle \mathbf{X}, \mathbf{D}, \mathbf{F} \rangle$, where $\mathbf{X}$ and $\mathbf{D}$ are the set of variables and their domains as defined in a CSP, and $\mathbf{F}$ is a set of *weighted constraints*. A weighted constraint $f_i \in \mathbf{F}$ is a function $f_i : \times_{x_j \in \mathbf{x}^i} D_j \rightarrow \mathbb{R}^+ \cup \{\perp\}$, where $\mathbf{x}^i \subseteq \mathbf{X}$ is the scope of f_i and $\perp$ is a special element used to denote that a given combination of values for the variables in $\mathbf{x}^i$ is not allowed, and it has the property that $a + \perp = \perp + a = \perp$, for all $a \in \mathbb{R}^+$. The *cost* of an assignment σ is the sum of the evaluation of the constraints involving all the variables in σ. A *solution* is a complete assignment with cost different from $\perp$, and an optimal solution is a solution with minimal cost.

Thus, a WCSP is a generalization of a CSP which, in turn, can be seen as a WCSP whose constraints use exclusively the costs 0 and $\perp$. The terms WCSP and *Constraint Optimization Problem*

---

1. The presence of a fixed ordering of variables is assumed.
2. A constraint with $k = 3$ is also called a *ternary* constraint and a constraint with $k = n$ is also called a *global* constraint.
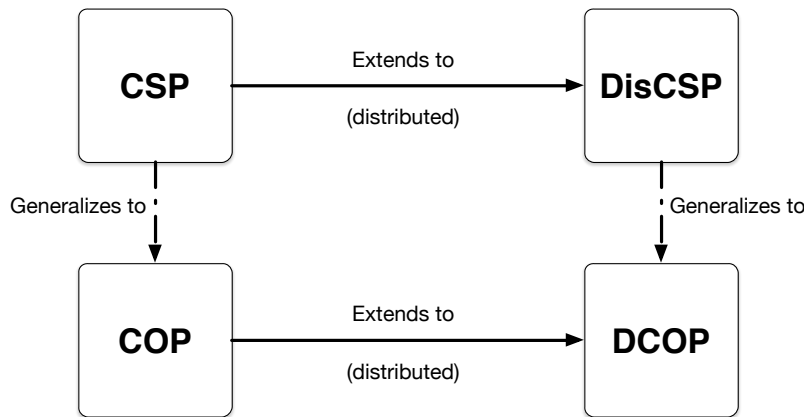
Figure 2: DCOP Problems as a Generalization and Extension of Constraint Satisfaction Problems

(COP) have been used interchangeably in the literature and the use of the latter term has been widely adopted in the recent years.

2.3 Distributed Constraint Satisfaction Problems

When the elements of a CSP are distributed among a set of autonomous agents, the resulting model is referred to as a *Distributed Constraint Satisfaction Problem (DisCSP)* (Yokoo, Durfee, Ishida, & Kuwabara, 1998; Yokoo, 2001). A DisCSP is a tuple $\langle \mathbf{A}, \mathbf{X}, \mathbf{D}, \mathbf{C}, \alpha \rangle$, where $\mathbf{X}$, $\mathbf{D}$, and $\mathbf{C}$ are the set of variables, their domains, and the set of constraints, as defined in a CSP; $\mathbf{A} = \{a_1, \ldots, a_m\}$ is a finite set of autonomous agents; and $\alpha : \mathbf{X} \to \mathbf{A}$ is a surjective function, from variables to agents, which assigns the control of each variable $x \in \mathbf{X}$ to an agent $\alpha(x)$. The goal in a DisCSP is to find a complete assignment that satisfies all the constraints of the problem.

DisCSPs can be seen as an extension of CSPs to the multi-agent case, where agents communicate with each other to assign values to the variables they control so as to satisfy all the problem constraints. For a survey on the topic, the interested reader is referred to the work of Rossi et al. (2006, ch. 20).

2.4 Distributed Constraint Optimization Problems

Similar to the generalization of CSPs to COPs, the *Distributed Constraint Optimization Problem (DCOP)* model (Modi et al., 2005; Petcu & Faltings, 2005b; Gershman et al., 2009; Yeoh & Yokoo, 2012) emerges as a generalization of the DisCSP model, where constraints specify a degree of preference over their violation, rather than a Boolean satisfaction metric. DCOPs can also be viewed as an extension of the COP framework to the multi-agent case, where agents control variables and constraints, and need to coordinate the value assignment for the variables they control so as to optimize a global objective function. The DCOP framework is formally introduced in the next section.

ELEMENT		CHARACTERIZATION	
AGENT(S)	BEHAVIOR	Deterministic	Stochastic
	KNOWLEDGE	Total	Partial
	TEAMWORK	Cooperative	Competitive
ENVIRONMENT	BEHAVIOR	Deterministic	Stochastic
	EVOLUTION	Static	Dynamic

Table 2: DCOP Classification Elements

3. DCOP Classification

The DCOP model has undergone a process of continuous evolution to capture diverse characteristics of the agent behavior and the environment in which agents operate. This section proposes a classification of DCOP models from a multi-agent systems perspective. It accounts for the different assumptions made about the behavior of the agents and their interactions with the environment. The classification is based on the following elements (summarized in Table 2):

- **Agent Behavior:** This parameter captures the stochastic nature of the *effects* of an action being executed. These effects can be either *deterministic* or *stochastic*.

- **Agent Knowledge:** This parameter captures the *knowledge* of an agent about its own state and the environment. It can be *total* or *partial* (i.e., *incomplete*).

- **Agent Teamwork:** This parameter characterizes the approach undertaken by (teams of) agents to solve a distributed problem. It can be either a *cooperative* or a *competitive* resolution approach. In the former class, all agents cooperate to achieve a common goal (i.e., they all optimize a global objective function). In the latter class, each agent (or team of agents) seeks to achieve its own individual goal (i.e., each agent optimizes its individual objective functions).

- **Environment Behavior:** This parameter captures the exogenous properties of the environment. The response of the environment to the execution of an action can be either *deterministic* or *stochastic*.

- **Environment Evolution:** This parameter captures whether the DCOP does not change over time (*static*) or it changes over time (*dynamic*).

Figure 3 illustrates a categorization of the DCOP models proposed to date from a MAS perspective. This survey focuses on the DCOP models proposed at the junction of constraint programming, game theory, and decision theory. The *classical* DCOP model is directly inherited from constraint programming as it extends the WCSP model to a distributed setting. It is characterized by a static model, a deterministic environment and agent behavior, a total agent knowledge, and a cooperative agent teamwork. Game theoretical concepts explored in the context of auctions and negotiations have influenced the DCOP framework leading to the development of the Asymmetric DCOP and the Multi-Objective DCOP. The DCOP framework has also borrowed fundamental decision theoretical concepts related to modeling uncertain and dynamic environments, resulting in models like the Probabilistic DCOP and the Dynamic DCOP. Researchers from the DCOP community have also designed solutions that inherit from all of the three communities.

The next sections describe the different DCOP frameworks, starting with classical DCOPs before proceeding to its various extensions. The survey focuses on a categorization based on three dimensions: *Agent knowledge*, *environment behavior*, and *environment evolution*. It assumes a *de-*

	Environment Behavior	
	DETERMINISTIC	STOCHASTIC
STATIC	Classical DCOP	Probabilistic DCOP
DYNAMIC	Dynamic DCOP	—

(Environment Evolution)
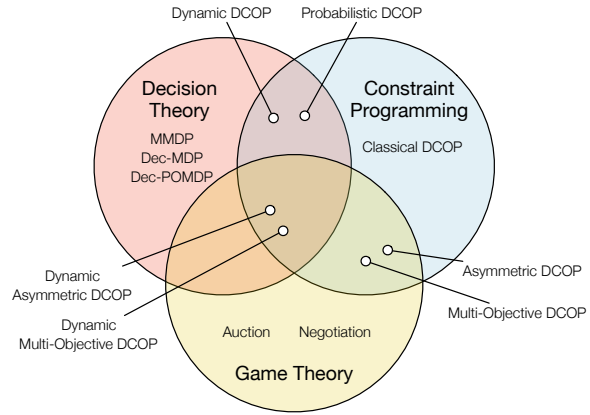

Table 3: DCOPs Models   Figure 3: DCOPs Within a MAS Perspective

terministic agent behavior, a *fully cooperative agent teamwork*, and a *total agent knowledge* (unless otherwise specified), as they are, by far, common assumptions adopted by the DCOP community. The DCOP models associated to this categorization are summarized in Table 3. The bottom-right entry of the table is left empty, indicating a promising model with dynamic and uncertain environments that, to the best of our knowledge, has not been explored yet. There has been only a modest amount of effort in modeling the different aspects of teamwork within the DCOP community. Section 7 describes a formalism that has been adopted to model DCOPs with mixed cooperative and competitive agents.

4. Classical DCOP

With respect to the proposed categorization, in the *classical DCOP* model (Modi et al., 2005; Petcu & Faltings, 2005b; Gershman et al., 2009; Yeoh & Yokoo, 2012) the agents are fully cooperative and have deterministic behavior and total knowledge. Additionally, the environment is static and deterministic. This section reviews the formal definitions of classical DCOPs, presents some relevant solving algorithms, and provides details of selected variants of classical DCOPs of particular interest.

4.1 Definition

A classical DCOP is described by a tuple $P = \langle \mathbf{A}, \mathbf{X}, \mathbf{D}, \mathbf{F}, \alpha \rangle$, where:
- $\mathbf{A} = \{a_1, \ldots, a_m\}$ is a finite set of agents.
- $\mathbf{X} = \{x_1, \ldots, x_n\}$ is a finite set of variables, with $n \geq m$.
- $\mathbf{D} = \{D_1, \ldots, D_n\}$ is a set of finite domains for the variables in $\mathbf{X}$, with D_i being the domain of variable x_i.
- $\mathbf{F} = \{f_1, \ldots, f_k\}$ is a finite set of *cost functions*, with $f_i : \times_{x_j \in \mathbf{x}^i} D_j \rightarrow \mathbb{R}^+ \cup \{\bot\}$, where similar to WCSPs, $\mathbf{x}^i \subseteq \mathbf{X}$ is the set of variables relevant to f_i, referred to as the *scope* of f_i. The *arity* of a cost function is the number of variables in its scope. Each cost function f_i represents a factor in a *global objective function* $\mathbf{F}_g(\mathbf{X}) = \sum_{i=1}^{k} f_i(\mathbf{x}^i)$. In the DCOP literature, the cost functions f_i are also called *constraints*, *utility functions*, or *reward functions*.

- $\alpha : \mathbf{X} \to \mathbf{A}$ is a total and onto function, from variables to agents, which assigns the control of each variable $x \in \mathbf{X}$ to an agent $\alpha(x)$.

With a slight abuse of notation, $\alpha(f_i)$ will be used to denote the set of agents whose variables are involved in the scope of f_i, i.e., $\alpha(f_i) = \{\alpha(x) \mid x \in \mathbf{x}^i\}$. A *partial* assignment is a value assignment for a proper subset of variables of $\mathbf{X}$. An assignment is *complete* if it assigns a value to each variable in $\mathbf{X}$. For a given complete assignment σ, we say that a cost function f_i is satisfied by σ if $f_i(\sigma_{\mathbf{x}^i}) \neq \bot$. A complete assignment is a *solution* of a DCOP if it satisfies all its cost functions. The goal in a DCOP is to find a solution that minimizes the total problem cost expressed by its cost functions:[3]

$$\sigma^* := \operatorname*{argmin}_{\sigma \in \Sigma} \mathbf{F}_g(\sigma) = \operatorname*{argmin}_{\sigma \in \Sigma} \sum_{f_i \in \mathbf{F}} f_i(\sigma_{\mathbf{x}^i}), \qquad (1)$$

where Σ is the *state space*, defined as the set of all possible solutions.

Given an agent a_i, $L_{a_i} = \{x_j \in \mathbf{X} \mid \alpha(x_j) = a_i\}$ denotes the set of variables controlled by agent a_i, or its *local variables*, and $N_{a_i} = \{a_i' \in \mathbf{A} \mid a_i \neq a_i', \exists f_j \in \mathbf{F}, x_r, x_s \in \mathbf{x}^j, \alpha(x_r) = a_i \wedge \alpha(x_s) = a_i'\}$ denotes the set of its *neighboring agents*. A cost function f_i is said to be *hard* if $\forall \sigma \in \Sigma$ we have that $f_i(\sigma_{\mathbf{x}^i}) \in \{0, \bot\}$. Otherwise, the cost function is said to be *soft*.

Finding an optimal solution for a classical DCOP is known to be NP-hard (Modi et al., 2005).

4.2 DCOP: Representation and Coordination

Representation in DCOPs plays a fundamental role, both from an agent coordination perspective and from an algorithmic perspective. This section discusses the most predominant representations adopted in various DCOP algorithms. It starts by describing some widely adopted assumptions regarding agent knowledge and coordination, which will apply throughout this document, unless otherwise stated:

(i) A variable and its domain are known exclusively to the agent controlling it and its neighboring agents.

(ii) Each agent knows the values of the cost function involving at least one of its local variables. No other agent has knowledge about such cost function.

(iii) Each agent knows (and it may communicate with) exclusively its own neighboring agents.

4.2.1 CONSTRAINT GRAPH

Given a DCOP P, $G_P = (\mathbf{X}, E_C)$ is the *constraint graph* of P, where an undirected edge $\{x, y\} \in E_C$ exists if and only if there exists $f_j \in \mathbf{F}$ such that $\{x, y\} \subseteq \mathbf{x}^j$. A constraint graph is a standard way to visualize a DCOP instance. It underlines the agents' locality of interactions and therefore it is commonly adopted by DCOP resolution algorithms.

Given an ordering o on $\mathbf{X}$, a variable x_i is said to have a higher *priority* with respect to a variable x_j if x_i appears before x_j in o. Given a constraint graph G_P and an ordering o on its nodes, the *induced graph* G_P^* on o is the graph obtained by connecting nodes, processed in increasing order of priority, to all their higher-priority neighbors. For a given node, the number of higher-priority

---

3. Alternatively, one can define a maximization problem by substituting the $\mathtt{argmin}$ operator in Equation 1 with $\mathtt{argmax}$. Typically, if the objective functions are referred to as utility functions or reward functions, then the DCOP is a maximization problem. Conversely, if the objective functions are referred to as cost functions, then the DCOP is a minimization problem.

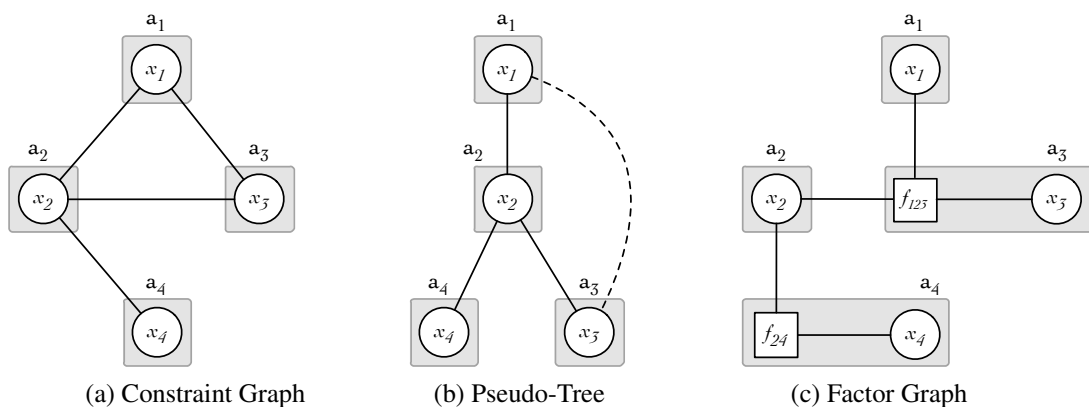(a) Constraint Graph (b) Pseudo-Tree (c) Factor Graph

Figure 4: DCOP Representations

neighbors is referred to as its *width*. The *induced width* w_o^* of G_P is the maximum width over all the nodes of G_P^* on ordering o.

Figure 4(a) shows an example constraint graph of a DCOP with four agents a_1 through a_4, each controlling one variable with domain {0,1}. There are two cost functions: a k-ary cost function f_{123} with scope $\mathbf{x}^{123} = \{x_1, x_2, x_3\}$ and represented by a clique among x_1, x_2, and x_3; and a binary cost function f_{24} with scope $\mathbf{x}^{24} = \{x_2, x_4\}$.

4.2.2 PSEUDO-TREE

A number of DCOP algorithms require a partial ordering among the agents. In particular, when such an order is derived from a depth-first search (DFS) exploration, the resulting structure is known as a *(DFS) pseudo-tree*. A pseudo-tree arrangement for a DCOP P is a subgraph $T_P = \langle \mathbf{X}, E_T \rangle$ of G_P such that T_P is a spanning tree of G_P – i.e., a connected subgraph of G_P containing all the nodes and being a rooted tree – with the following additional condition: for each $x, y \in \mathbf{X}$, if $\{x, y\} \subseteq \mathbf{x}^i$ for some $f_i \in \mathbf{F}$, then x, y appear in the same branch of T_P (i.e., x is an ancestor of y in T_P or vice versa). Edges of G_P that are *in* (respectively *out* of) E_T are called *tree edges* (respectively *backedges*). The tree edges connect parent-child nodes, while backedges connect a node with its *pseudo-parents* and its *pseudo-children*. The *separator* of an agent a_i is the set containing all the ancestors of a_i in the pseudo-tree (through tree edges or backedges) that are connected to a_i or to one of its descendants. The notation C_{a_i}, PC_{a_i}, P_{a_i}, and PP_{a_i} will be used to indicate the set of children, pseudo-children, parent, and pseudo-parents of the agent a_i.

Both constraint graph and pseudo-tree representations cannot deal explicitly with k-ary cost functions (with $k > 2$). A typical artifact to deal with such cost functions in a pseudo-tree representation is to introduce a virtual variable that monitors the value assignments for all the variables in the scope of the cost function, and generates the cost values (Bowring, Tambe, & Yokoo, 2006) – the role of the virtual variables can be delegated to one of the variables participating in the cost function (Pecora, Modi, & Scerri, 2006; Matsui, Matsuo, Silaghi, Hirayama, & Yokoo, 2008).

Figure 4(b) shows one possible pseudo-tree of the example DCOP in Figure 4(a), where $C_{a_1} = \{x_2\}$, $PC_{a_1} = \{x_3\}$, $P_{a_4} = \{x_2\}$, and $PP_{a_3} = \{x_1\}$. The solid lines are tree edges and dotted lines are backedges.

4.2.3 Factor Graph

Another way to represent DCOPs is through a *factor graph* (Kschischang, Frey, & Loeliger, 2001). A factor graph is a bipartite graph used to represent the factorization of a function. In particular, given the global objective function $\mathbf{F}_g$, the corresponding factor graph $F_P = \langle \mathbf{X}, \mathbf{F}, E_F \rangle$ is composed of variable nodes $x_i \in \mathbf{X}$, factor nodes $f_j \in \mathbf{F}$, and edges E_F such that there is an undirected edge between factor node f_j and variable node x_i if $x_i \in \mathbf{x}^j$.

Factor graphs can handle k-ary cost functions explicitly. To do so, they use a similar method as the one adopted within pseudo-trees with such cost functions: They delegate the control of a factor node to one of the agents controlling a variable in the scope of the cost function. From an algorithmic perspective, the algorithms designed over factor graphs can directly handle k-ary cost functions, while algorithms designed over pseudo-trees require changes in the algorithm design so to delegate the control of the k-ary cost functions to some particular entity.

Figure 4(c) shows the factor graph of the example DCOP in Figure 4(a), where each agent a_i controls its variable x_i and, in addition, a_3 controls the cost function f_{123} and a_4 controls cost function f_{24}.

4.3 Algorithms

The field of classical DCOPs is mature and a number of different resolution algorithms have been proposed. DCOP algorithms can be classified as being either *complete* or *incomplete*, based on whether they can guarantee the optimal solution or they trade optimality for shorter execution times, producing near-optimal solutions. They can also be characterized based on their runtime characteristics, their memory requirements, and their communication requirements (e.g., the number and size of messages that they send and whether they communicate with their neighboring agents only or also to non-neighboring agents). Table 4 tabulates the properties of a number of key DCOP algorithms that will be surveyed in Sections 4.3.4 and 4.3.5. An algorithm is said *anytime* if it can return a valid solution even if the DCOP agents are interrupted at any time before the algorithm terminates. Anytime algorithms are expected to seek for solutions of increasing quality as they keep running (Zivan, Okamoto, & Peled, 2014).

All these algorithms were originally developed under the assumption that each agent controls exactly one variable. The description of their properties will follow the same assumption. These properties may change when generalizing the algorithms to allow for agents to control multiple variables, but they will depend on how the algorithms are generalized. Throughout this document, the following notation will be often adopted when discussing the complexity of the algorithms:

- $n = |\mathbf{X}|$ refers to the number of variables in the problem; in Table 4, n also refers to the number of agents in the problem since each agent has exactly one variable;
- $d = \max_{D_i \in \mathbf{D}} |D_i|$ refers to the size of the largest domain;
- w^* refers to the induced width of the pseudo-tree;
- $l = \max_{a_i \in \mathbf{A}} |N_{a_i}|$ refers to the largest number of neighboring agents; and
- ℓ refers to the number of iterations in incomplete algorithms.

In addition, each of these classes can be categorized into several groups, depending on the degree of locality exploited by the algorithms, the way local information is updated, and the type of exploration process adopted. These different categories are described next.

Algorithm	Quality Characteristics		Runtime Characteristics		Memory	Communication Characteristics		
	Optimal?	Error Bound?	Complexity	Anytime?	per Agent	# Messages	Message Size	Local Communication?
SyncBB	✓	✓	$O(d^n)$	✓	$O(n)$	$O(d^n)$	$O(n)$	×
AFB	✓	✓	$O(d^n)$	✓	$O(n)$	$O(d^n)$	$O(n)$	×
ADOPT	✓	✓	$O(d^n)$	×	$O(n+ld)$	$O(d^n)$	$O(n)$	✓
ConcFB	✓	✓	$O(d^n)$	✓	$O(n)$	$O(d^n)$	$O(n)$	×
DPOP	✓	✓	$O(d^{w^*})$	×	$O(d^{w^*})$	$O(n)$	$O(d^{w^*})$	✓
OptAPO	✓	✓	$O(d^n)$	×	$O(ld)$	$O(d^n)$	$O(d+n)$	×
Max-Sum	×	×	$O(\ell d^l)$	×	$O(d^l)$	$O(\ell nl)$	$O(d)$	✓
Region Optimal	×	✓	$O(\ell d^{w^*})$	✓	$O(d^{w^*})$	$O(\ell n^2)$	$O(d^{w^*})$	×
MGM	×	×	$O(\ell ld)$	✓	$O(l)$	$O(\ell nl)$	$O(1)$	✓
DSA	×	×	$O(\ell ld)$	✓	$O(l)$	$O(\ell nl)$	$O(1)$	✓
DUCT	×	✓	$O(\ell ld)$	✓	$O(d^{w^*})$	$O(\ell n)$	$O(n)$	✓
D-Gibbs	×	✓	$O(\ell ld)$	✓	$O(l)$	$O(\ell nl)$	$O(1)$	✓

Table 4: Quality, Runtime, Memory, and Communication Characteristics of DCOP Algorithms

4.3.1 PARTIAL CENTRALIZATION

In general, the DCOP solving process is decentralized, driving DCOP algorithms to follow the agent knowledge and communication restrictions described in Section 4.2. However, some algorithms explore methods to centralize the decisions to be taken by a group of agents, by delegating them to one of the agents in the group. These algorithms explore the concept of *partial centralization* (Hirayama & Yokoo, 1997; Mailler & Lesser, 2004; Petcu, Faltings, & Mailler, 2007), and thus they are classified as *partially centralized* algorithms. Typically, partial centralization improves the algorithms' performance allowing agents to coordinate their local assignments more efficiently. However, such performance enhancement comes with a loss of information privacy, as the centralizing agent needs to be granted access to the local subproblem of other agents in the group (Greenstadt, Grosz, & Smith, 2007; Mailler & Lesser, 2004). In contrast, *fully* decentralized algorithms inherently reduce the amount of information privacy at cost of a larger communication effort.

4.3.2 SYNCHRONICITY

DCOP algorithms can enhance their effectiveness by exploiting distributed and parallel processing. Based on the way the agents update their local information, DCOP algorithms are classified as *synchronous* or *asynchronous*. Asynchronous algorithms allow agents to update the assignment for their variables based solely on their local view of the problem, and thus independently from the actual decisions of the other agents (Modi et al., 2005; Farinelli, Rogers, Petcu, & Jennings, 2008; Gershman et al., 2009). In contrast, synchronous algorithms constrain the agents decisions to follow a particular order, typically enforced by the representation structure adopted (Mailler & Lesser, 2004; Petcu & Faltings, 2005b; Pearce & Tambe, 2007).

Synchronous algorithms tend to delay the actions of some agents guaranteeing that their local view of the problem is always consistent with that of the other agents. In contrast, asynchronous algorithms tend to minimize the idle-time of the agents, which in turn can react quickly to each message being processed; however, they provide no guarantee on the consistency of the state of the local view of each agent. Such effect has been studied by Peri and Meisels (2013), concluding that inconsistent agents' views may cause a negative impact on network load and algorithm performance, and that introducing some level of synchronization may be beneficial for some algorithms, enhancing their performance.
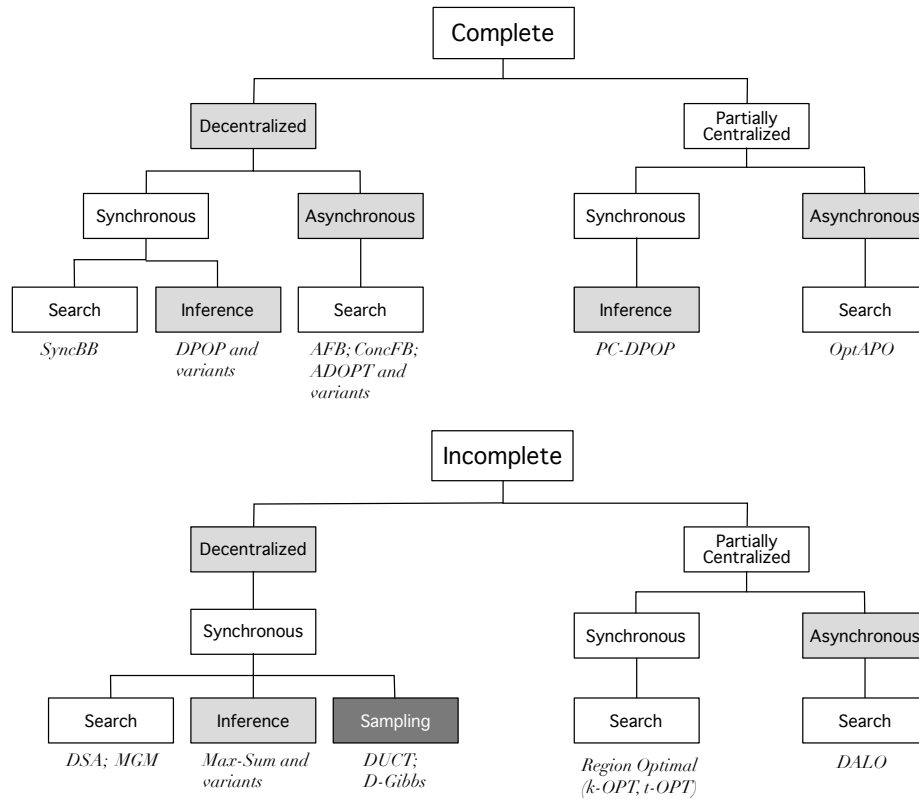
Figure 5: Classical DCOP Algorithm Taxonomy

4.3.3 EXPLORATION PROCESS

The resolution process adopted by each algorithm can be classified in three categories (Yeoh, 2010):

- *Search-based algorithms* are based on the use of search techniques to explore the space of possible solutions. These algorithms are often derived from corresponding search techniques developed for centralized AI search problems, such as best-first search and depth-first search.
- *Inference-based algorithms* are derived from dynamic programming and belief propagation techniques. These algorithms allow agents to exploit the structure of the constraint graph to aggregate costs from their neighbors, effectively reducing the problem size at each step of the algorithm.
- *Sampling-based algorithms* are incomplete approaches that sample the search space to approximate a function (typically, a probability distribution) as a product of statistical inference.

Figure 5 illustrates a taxonomy of classical DCOP algorithms. The following subsections summarize some representative complete and incomplete algorithms from each of the classes introduced above. A detailed description of the DCOP algorithms is beyond the scope of this manuscript. The interested reader is referred to the original articles that introduce each algorithm.

4.3.4 COMPLETE ALGORITHMS

Some of the algorithms described below were originally designed to solve the variant of DCOPs that maximizes rewards, while others solve the variant that minimizes costs. However, the algo-

rithms that maximize rewards can be easily adapted to minimize costs. For consistency, this survey describes the version of the algorithms that focus on minimization of costs. It also describes their quality, runtime, memory, and communication characteristics as summarized in Table 4.

SyncBB (Hirayama & Yokoo, 1997). *Synchronous Branch-and-Bound (SyncBB)* is a complete, synchronous, search-based algorithm that can be considered as a distributed version of a branch-and-bound algorithm. It uses a complete ordering of the agents to extend a *Current Partial Assignment (CPA)* via a synchronous communication process. The CPA holds the assignments of all the variables controlled by all the visited agents, and, in addition, functions as a mechanism to propagate bound information. The algorithm prunes those parts of the search space whose solution quality is sub-optimal, by exploiting the bounds that are updated at each step of the algorithm.

SyncBB agents perform $O(d^n)$ number of operations since the lowest priority agent needs to enumerate through all possible value combinations for all variables. While, by default, it is not an anytime algorithm, it can be easily extended to have an anytime property since it is a branch-and-bound algorithm. The memory requirement per SyncBB agent is $O(n)$ since the lowest priority agent stores the value assignment of all problem variables. In terms of communication requirement, SyncBB agents send $O(d^n)$ number of messages: The lowest priority agent enumerates through all possible value combinations for all variables and sends a message for each combination. The largest message, which contains the value assignment of all variables, is of size $O(n)$. Finally, the communication model of SyncBB depends on the given agent's complete ordering. Thus, agents may communicate with non-neighboring agents.

AFB (Gershman et al., 2009). *Asynchronous Forward Bounding (AFB)* is a complete, asynchronous, search-based algorithm. It can be considered as an asynchronous version of SyncBB. In this algorithm, agents communicate their cost estimates, which in turn are used to compute bounds and prune the search space. In AFB, agents extend a CPA sequentially, provided that the lower bound on their costs does not exceed the global upper bound, that is, the cost of the best solution found so far. Each agent performing an assignment (the "assigning" agent) triggers asynchronous checks of bounds, by sending *forward* messages containing copies of the CPA to agents that have not yet assigned their variables. The unassigned agents that receive a CPA estimate the lower bound of the CPA given their local view of the constraint graph and send their estimates back to the agent that originated the *forward* message. This assigning agent will receive these estimates asynchronously and aggregate them into an updated lower bound. If the updated lower bound exceeds the current upper bound, the agent initiates a backtracking phase.

The runtime, memory, and communication characteristics of AFB are identical to those of SyncBB for the same reasons. However, while both AFB and SyncBB agents communicate with non-neighboring agents, AFB agents broadcasts some of their messages while SyncBB agents do not.

ADOPT (Modi et al., 2005). *Asynchronous Distributed OPTimization (ADOPT)* is a complete, asynchronous, search-based algorithm. It can be considered as a distributed version of a memory-bounded best-first search algorithm. It makes use of a DFS pseudo-tree ordering of the agents. The algorithm relies on maintaining, in each agent, lower and upper bounds on the solution cost for the subtree rooted at its node in the DFS tree. Agents explore partial assignments in best-first order, that is, in increasing lower bound order. They use *COST* messages (propagated upwards in the DFS pseudo-tree) and *THRESHOLD* and *VALUE* messages (propagated downwards in the pseudo-tree) to iteratively tighten the lower and upper bounds, until the lower bound of the minimum cost

solution is equal to its upper bound. ADOPT agents store lower bounds as thresholds, which can be used to prune partial assignments that are provably sub-optimal.

Similar to SyncBB and AFB, ADOPT agents perform $O(d^n)$ number of operations since the lowest priority agent needs to enumerate through all possible value combinations for all variables when the pseudo-tree degenerates into a pseudo-chain. It is also not an anytime algorithm as it is a best-first search algorithm. The memory requirement per ADOPT agent is $O(n + ld)$, where $O(n)$ is used to store a *context*, which is the value assignment of all higher-priority variables, and $O(ld)$ is used to store the lower and upper bounds for each domain value and variable belonging to the agent's child agents. Finally, ADOPT agents communicate exclusively with their neighboring agents.

ADOPT has been extended in several ways. In particular, *BnB-ADOPT* (Yeoh, Felner, & Koenig, 2010; Gutierrez & Meseguer, 2012b) uses a branch-and-bound method to reduce the amount of computation performed during search, and *ADOPT(k)* combines both ADOPT and BnB-ADOPT into an integrated algorithm (Gutierrez, Meseguer, & Yeoh, 2011). There are also extensions that trade solution optimality for smaller runtimes (Yeoh, Sun, & Koenig, 2009a), extensions that use more memory for smaller runtimes (Yeoh, Varakantham, & Koenig, 2009b), and extensions that maintain soft arc-consistency (Bessiere, Gutierrez, & Meseguer, 2012; Bessiere, Brito, Gutierrez, & Meseguer, 2014; Gutierrez & Meseguer, 2012a; Gutierrez, Lee, Lei, Mak, & Meseguer, 2013).

Finally, the *No-Commitment Branch and Bound (NCBB)* algorithm (Chechetka & Sycara, 2006) can be considered as a variant of ADOPT and SyncBB. Similar to ADOPT, NCBB agents exploit the structure defined by a pseudo-tree order to decompose the global objective function. This allow the agents to search non-intersecting parts of the search space concurrently. Another main feature of NCBB is the eager propagation of lower bounds on solution cost: An NCBB agent propagates its lower bound every time it learns about its ancestors' assignments. This feature provides an efficient pruning of the search space. The runtime, memory, and communication characteristics of NCBB are the same as those of ADOPT except that NCBB is an anytime algorithm.

ConcFB (Netzer, Grubshtein, & Meisels, 2012). *Concurrent Forward Bounding (ConcFB)* is a complete, asynchronous, search-based algorithm that runs multiple parallel versions of AFB concurrently. By running multiple concurrent search procedures, it is able to quickly find a solution, apply a forward bounding process to detect regions of the search space to prune, and to dynamically create new search processes when detecting promising sub-spaces. Similar to AFB, it uses a complete ordering of agents and variables instead of pseudo-trees. As such, it is able to simplify the management of reordering heuristics, which can provide substantial speed up to the search process (Zivan & Meisels, 2006).

The algorithm operates as follows: Each agent maintains a global upper bound, which is updated during the search process. The highest-priority agent begins the process by generating a number of different search processes (SP), one for each value of its variable. It then sends an *LB_Request* message to all unassigned agents. This *LB_Request* message contains the current CPA and triggers a calculation of the lower bounds of the receiving agents, which are sent back to the sender agent via a *LB_Report* message. If the sum of the aggregated costs and the current CPA cost is no smaller than the current upper bound, the agent selects another value for its variable and repeats the process. If the agent has exhausted all value assignments for its variable, then it backtracks, sending the CPA to the last assigning agent. If the CPA cost is lower than the current upper bound, then it forwards the *CPA* message to the next non-assigned agent. Upon receiving a *CPA* message, the agent repeats

the above process. When the lowest-priority agent finds a solution resulting to a new upper bound, it broadcasts the upper bound via a *UB* message, which is stored by each each agent.

Netzer et al. (2012) described a series of enhancements that can be used to speed up the search process of ConcFB, including dynamic variable ordering and dynamic splitting. Despite the process within a subproblem is carried out in a synchronous fashion, different subproblems are explored independently. Thus, the agents act asynchronously and concurrently. The runtime, memory, and communication characteristics of ConcFB are identical to those of AFB since it runs multiple parallel versions of AFB concurrently.

DPOP (Petcu & Faltings, 2005b). *Distributed Pseudo-tree Optimization Procedure (DPOP)* is a complete, synchronous, inference-based algorithm that makes use of a DFS pseudo-tree ordering of the agents. It involves three phases. In the first phase, the agents order themselves into a DFS pseudo-tree. In the second phase, called the UTIL propagation phase, each agent, starting from the leaves of the pseudo-tree, aggregates the costs in its subtree for each value combination of variables in its separator. The aggregated costs are encoded in a *UTIL* message, which is propagated from children to their parents, up to the root. In the third phase, called the VALUE propagation phase, each agent, starting from the root of the pseudo-tree, selects the optimal value for its variable. The optimal values are calculated based on the UTIL messages received from the agent's children and the *VALUE* message received from its parent. The *VALUE* messages contain the optimal values of the agents and are propagated from parents to their children, down to the leaves of the pseudo-tree.

DPOP agents perform $O(d^{w^*})$ number of operations. When an agent optimizes for each value combination of variables in its separator, it takes $O(d^{w^*})$ operations since there are w^* variables in the separator set in the worst case. It is not an anytime algorithm as it terminates upon finding its first solution, which is an optimal solution. The memory requirement per DPOP agent is $O(d^{w^*})$ since it stores all value combinations of variables in its separator. In terms of communication requirement, DPOP agents send $O(n)$ messages in total; $O(n)$ UTIL messages are propagated up the pseudo-tree and $O(n)$ VALUE messages are propagated down the pseudo-tree. The largest message sent by an agent, which contains the aggregated costs in its subtree for each value combination of variables in its separator, is $O(d^{w^*})$. Finally, DPOP agents only communicate with their neighboring agents only.

DPOP has also been extended in several ways to enhance its performance and capabilities. O-DPOP and MB-DPOP trade runtimes for smaller memory requirements (Petcu & Faltings, 2006, 2007a), A-DPOP trades solution optimality for smaller runtimes (Petcu & Faltings, 2005a), SS-DPOP trades runtime for increased privacy (Greenstadt et al., 2007), PC-DPOP trades privacy for smaller runtimes (Petcu et al., 2007), H-DPOP propagates hard constraints for smaller runtimes (Kumar, Petcu, & Faltings, 2008), BrC-DPOP enforces branch consistency for smaller runtimes (Fioretto, Le, Yeoh, Pontelli, & Son, 2014b), and ASP-DPOP is a declarative version of DPOP that uses Answer Set Programming (Le, Son, Pontelli, & Yeoh, 2015).

OptAPO (Mailler & Lesser, 2004). *Optimal Asynchronous Partial Overlay (OptAPO)* is a complete, asynchronous, search-based algorithm. It trades agent privacy for smaller runtimes through partial centralization. It employs a cooperative mediation schema, where agents can act as mediators and propose value assignments to other agents. In particular, the agents check if there is a conflicting assignment with some neighboring agent. If a conflict is found, the agent with the highest priority acts as a mediator. During mediation, OptAPO solves subproblems using a centralized branch-and-bound-based search, and when solutions of overlapping subproblems still have

conflicting assignments, the solving agents increase the degree of centralization to resolve them. By sharing their knowledge with centralized entities, agents can improve their local decisions, reducing the communication costs. For instance, the algorithm has been shown to be superior to ADOPT on simple combinatorial problems (Mailler & Lesser, 2004). However, it is possible that several mediators solve overlapping problems, duplicating efforts (Petcu et al., 2007), which can be a bottleneck in dense problems.

OptAPO agents perform $O(d^n)$ number of operations, in the worst case, as a mediator agent may solve the entire problem. Like ADOPT and DPOP, OptAPO is not an anytime algorithm. The memory requirement per OptAPO agent is $O(nd)$ since it needs to store all value combinations of variables in its mediation group, which is of size $O(n)$. In terms of communication requirement, OptAPO agents send $O(d^n)$ messages in the worst case, though the number of messages decreases with increasing partial centralization. The size of the messages is bounded by $O(d + n)$, where in the initialization phase of each mediation step, each agent sends its domain to its neighbors and the list of variables that it seeks to mediate. Finally, OptAPO agents can communicate with non-neighboring agents during the mediation phase.

The original version of OptAPO has been shown to be incomplete due to the asynchronicity of the different mediators' groups, which can lead to race conditions. Grinshpoun and Meisels (2008) proposed a complete variant that remedies this issue.

4.3.5 Incomplete Algorithms

Max-Sum (Farinelli et al., 2008). *Max-Sum* is an incomplete, synchronous, inference-based algorithm based on belief propagation. It operates on factor graphs by performing a marginalization process of the cost functions, and optimizing the costs for each given variable. This process is performed by recursively propagating messages between variable nodes and factor nodes. The value assignments take into account their impact on the marginalized cost function. Max-Sum is guaranteed to converge to an optimal solution in acyclic graphs, but convergence is not guaranteed on cyclic graphs.

Max-Sum agents perform $O(d^l)$ number of operations in each iteration, where each agent needs to optimize for all value combinations of neighboring variables. It is not an anytime algorithm. The memory requirement per Max-Sum agent is $O(d^l)$ since it needs to store all value combinations of neighboring variables. In terms of communication requirement, in the worst case, each Max-Sum agent sends $O(l)$ messages in each iteration, one to each of its neighbor. Thus, the total number of messages sent across all agents is $O(\ell n l)$. Each message is of size $O(d)$ as it needs to contain the current aggregated costs of all the agent's variable's values. Finally, the agents communicate exclusively with their neighboring agents.

Max-Sum has been extended in several ways. Bounded Max-Sum bounds the quality of the solutions found by removing a subset of edges from a cyclic DCOP graph to make it acyclic, and running Max-Sum to solve the acyclic problem (Rogers, Farinelli, Stranders, & Jennings, 2011); Improved Bounded Max-Sum improves on the error bounds (Rollon & Larrosa, 2012); and Max-Sum_ADVP guarantees convergence in acyclic graphs through a two-phase value propagation phase (Zivan & Peled, 2012; Chen, Deng, & Wu, 2017). Max-Sum and its extensions have been successfully used to solve a number of large scale, complex MAS applications (see Section 8).

Region Optimal (Pearce & Tambe, 2007). *Region-optimal* algorithms are incomplete, synchronous, search-based algorithms that allow users to specify regions of the constraint graph and solve the sub-

problem within each region optimally. Regions may be defined to have a maximum size of k agents (Pearce & Tambe, 2007), t hops from each agent (Kiekintveld, Yin, Kumar, & Tambe, 2010), or a combination of both size and hops (Vinyals, Shieh, Cerquides, Rodriguez-Aguilar, Yin, Tambe, & Bowring, 2011). The concept of k-optimality is defined with respect to the number of agents whose assignments conflict, whose set is denoted by $\mathbf{c}(\sigma, \sigma')$, for two assignments σ and σ'. The deviating cost of σ with respect to σ', denoted by $\Delta(\sigma, \sigma')$, is defined as the difference of the aggregated cost associated to the assignment σ ($F(\sigma)$) minus the cost associated to σ' ($F(\sigma')$). An assignment σ is k-optimal if $\forall \sigma' \in \Sigma$, such that $|\mathbf{c}(\sigma, \sigma')| \leq k$, we have that $\Delta(\sigma, \sigma') \geq 0$. In contrast, the concept of t-distance emphasizes the number of hops from a central agent a of the region $\Omega_t(a)$, that is the set of agents which are separated from a by at most t hops. An assignment σ is t-distance optimal if, $\forall \sigma' \in \Sigma$, $F(\sigma) \geq F(\sigma')$ with $\mathbf{c}(\sigma, \sigma') \subseteq \Omega_t(a)$, for any $a \in \mathbf{A}$. Therefore, the solutions found have theoretical error bounds that are a function of k and/or t. Region-optimal algorithms adopt a partially-centralized resolution scheme in which the subproblem within each region is solved optimally by a centralized authority (Tassa, Zivan, & Grinshpoun, 2016). However, this scheme can be altered to use a distributed algorithm to solve each subproblem.

Region-optimal agents perform $O(d^{w^*})$ number of operations in each iteration, as each agent runs DPOP to solve the problem within each region optimally. It is also an anytime algorithm as solutions of improving quality are found until they are region-optimal. The memory requirement per region-optimal agent is $O(d^{w^*})$ since its region may have an induced width of w^* and it uses DPOP to solve the problem within its region. In terms of communication requirement, each region-optimal agent sends $O(n)$ messages, one to each agent within its region. Thus, the total number of messages sent across all agents is $O(\ell n^2)$. Each message is of size $O(d^{w^*})$ as it uses DPOP. Finally, the agents communicate to all agents within their region – either to a distance of $\lfloor \frac{k}{2} \rfloor$ or t hops away. Thus, they may communicate with non-neighboring agents.

An asynchronous version of regional-optimal algorithms, called *Distributed Asynchronous Local Optimization* (DALO), was proposed by Kiekintveld et al. (2010). The DALO simulator provides a mechanism to coordinate the decision of local groups of agents based on the concepts of k-optimality and t-distance.

MGM (Maheswaran, Pearce, & Tambe, 2004a). *Maximum Gain Message (MGM)* is an incomplete, synchronous, search-based algorithm that performs a distributed local search. Each agent starts by assigning a random value to each of its variables. Then, it sends this information to all its neighbors. Upon receiving the values of its neighbors, it calculates the maximum gain (i.e., the maximum decrease in cost) if it changes its value and sends this information to all its neighbors. Upon receiving the gains of its neighbors, the agent changes its value if its gain is the largest among those of its neighbors. This process repeats until a termination condition is met. MGM provides no quality guarantees on the returned solution.

MGM agents perform $O(ld)$ number of operations in each iteration, as each agent needs to compute the cost for each of its values by taking into account the values of all its neighbors. MGM is anytime since agents only change their values when they have a non-negative gain. The memory requirement per MGM agent is $O(l)$. Each agent needs to store the values of all its neighboring agents. In terms of communication requirement, each MGM agent sends $O(l)$ messages, one to each of its neighboring agents. Thus, the total number of messages sent across all agents is $O(\ell nl)$. Each message is of constant size $O(1)$ as it contains either the agent's current value or the agent's current gain. Finally, the agents communicate exclusively with their neighboring agents.

DSA (Zhang, Wang, Xing, & Wittenberg, 2005). *Distributed Stochastic Algorithm (DSA)* is an incomplete, synchronous, search-based algorithm that is similar to MGM, except that each agent does not send its gains to its neighbors and it does not change its value to the value with the maximum gain. Instead, it decides stochastically if it takes on the value with the maximum gain or other values with smaller gains. This stochasticity allows DSA to escape from local minima. Similar to MGM, it repeats the process until a termination condition is met, and it cannot provide quality guarantees on the returned solution. The runtime, memory, and communication characteristics of DSA are identical to those of MGM since it is essentially a stochastic variant of MGM.

DUCT (Ottens, Dimitrakakis, & Faltings, 2017). The *Distributed Upper Confidence Tree (DUCT)* algorithm is an incomplete, synchronous, sampling-based algorithm that is inspired by Monte-Carlo Tree Search and employs confidence bounds to solve DCOPs. DUCT emulates a search process analogous to that of ADOPT, where agents select the values to assign to their variables according to the information encoded in their *context* messages (i.e., the assignments to all the variables in the receiving variable's separator). However, rather than systematically selecting the next value to assign to their own variables, DUCT agents sample such values. To focus on promising assignments, DUCT constructs a confidence bound B, such that cost associated to the best value for any context is at least B, and hence agents sample the choice with the lowest bound. This process is started by the root agent of the pseudo-tree: After sampling a value for its variable, it communicates its assignment to its children in a *context* message. When an agent receives this message, it repeats this process until the leaf agents are reached. When the leaf agents choose a value assignment, they calculate the cost within their context and propagate this information up to the tree in a *cost* message. This process continues for a given number of iterations or until convergence is achieved, i.e., until the sampled values in two successive iterations do not change. Therefore, DUCT is able to provide quality guarantees on the returned solution.

DUCT agents perform $O(ld)$ number of operations in each iteration, as each agent needs to compute the cost for each of its values by taking into account the values of all its neighbors. It is an anytime algorithm; The quality guarantee improves with increasing number of iterations. The memory requirement per DUCT agent is $O(d^n)$[4] since it needs to store the best cost for all possible contexts. In terms of communication requirement, in each iteration, each DUCT agent sends one message to its parent in the pseudo-tree and one message to each of its children in the pseudo-tree. Thus, the total number of messages sent across all agents is $O(\ell n)$. Each message is of size $O(n)$; *context* messages contain the value assignment for all higher priority agents. Finally, the agents communicate exclusively with their neighboring agents.

D-Gibbs (Nguyen, Yeoh, & Lau, 2013). The *Distributed Gibbs (D-Gibbs)* algorithm is an incomplete, synchronous, sampling-based algorithm that extends the Gibbs sampling process (Geman & Geman, 1984) by tailoring it to solve DCOPs in a decentralized manner. The Gibbs sampling process is a centralized Markov Chain Monte-Carlo algorithm that can be used to approximate joint probability distributions. By mapping DCOPs to maximum a-posteriori estimation problems, probabilistic inference algorithms like Gibbs sampling can be used to solve DCOPs.

Like DUCT, it too operates on a pseudo-tree, and the agents sample sequentially from the root of the pseudo-tree down to the leaves. Like DUCT, each agent also stores a context (i.e., the current assignment to all the variables in its separator) and it samples based on this information. Specifically,

---

4. It is actually $O(d^t)$, where t is the depth of the pseudo-tree. However, in the worst case, when the pseudo-tree degenerates into a pseudo-chain, then $t = n$.

it computes the probability for each of its values given its context and chooses its current value based on this probability distribution. After it chooses its value, it informs its lower priority neighbors of its value, and its children agents start to sample. This process continues until all the leaf agents sample. Cost information is propagated up the pseudo-tree. This process continues for a fixed number of iterations or until convergence. Like DUCT, D-Gibbs is also able to provide quality guarantees on the returned solution.

The runtime characteristics of D-Gibbs are identical to that of DUCT and for the same reasons. However, its memory requirements are smaller: The memory requirement per D-Gibbs agent is $O(l)$ since it needs to store the current values of all its neighbors. In terms of communication requirement, in each iteration, each D-Gibbs agent sends $O(l)$ messages, one to each of its neighbors. Thus, the total number of messages sent across all agents is $O(\ell n l)$. Each message is of constant size $O(1)$ since they contain only the current value of the agent or partial cost of its solution. Finally, the agents communicate exclusively with their neighboring agents.

A version of the algorithm that speeds up the agents' sampling process with Graphical Processing Units (GPUs) is described by Fioretto, Yeoh, and Pontelli (2016a).

4.4 Tradeoffs Between the Various DCOP Algorithms

The various DCOP algorithms discussed above provide a good coverage across various characteristics that may be important in different applications. As such, how well suited an algorithm is for an application depends on how well the algorithm's characteristics match up to the application's characteristics. The next section discusses several suggestions for the types of algorithms that are recommended based on the characteristics of the application at hand.

4.4.1 COMPLETE ALGORITHMS

When optimality is a requirement of the application, then one is limited to complete algorithms:

- If the agents in the application have large amounts of memory and it is faster to send few large messages than many small messages, then inference-based algorithms (e.g., DPOP and its extensions) are preferred over search-based algorithms (e.g., SyncBB, AFB, ADOPT, ConcFB, OptAPO). This is because, in general, search algorithms perform some amount of redundant communication. Thus, for a given problem instance, the overall runtime of inference-based algorithms tend to be smaller than the runtime of search-based ones.
- If the agents in the application have limited amounts of memory, then one has to use the search-based algorithms (e.g., SyncBB, AFB, ADOPT, ConcFB, OptAPO), which have small memory requirements. The exception is when the problem has a small induced width (e.g., the constraint graph is acyclic), in which case inference-based algorithms (e.g., DPOP) are also preferred.
  - If partial centralization is allowed by the application, then OptAPO is preferred as it has been shown to outperform many of the other search algorithms (Mailler & Lesser, 2004).
  - Otherwise, ConcFB is recommended as it has been shown to outperform AFB due to the concurrent search (Netzer et al., 2012), and AFB has been shown to outperform ADOPT and SyncBB (Gershman et al., 2009). The exception is if the application does not permit agents to communicate directly to non-neighbors, in which case ConcFB, AFB, and SyncBB cannot be used and one is restricted to use ADOPT or one of its variants. Note that many of the variants (e.g., BnB-ADOPT, NCBB) have been shown to significantly outperform ADOPT while main-

taining the same runtime, memory, and communication requirements (Chechetka & Sycara, 2006; Yeoh et al., 2010).

4.4.2 INCOMPLETE ALGORITHMS

In terms of incomplete algorithms, the following recommendations are given:

- If the solution returned must have an accompanying quality guarantee, then, one can choose to use Bounded Max-Sum, region-optimal algorithms, DUCT, or D-Gibbs. Bounded Max-Sum allows users to choose the error bound as a function of the different subsets of edges that can be removed from the graph to make it acyclic. Region-optimal algorithms allow users to parameterize the error bound according to the size of the region k or the number of hops t that the solution should be optimal for. Finally, DUCT and D-Gibbs allow users to parameterize the error bound based on the number of sampling iterations to conduct. The error bounds for these two algorithms are also probabilistic bounds (i.e., the likelihood that the quality of the solution is within an error bound is a function of the number of iterations). Therefore, the choice of algorithm will depend on the type of error bound one would like to impose on the solutions. One may also choose to use a number of extensions of complete algorithms (e.g., Weighted (BnB-)ADOPT and A-DPOP) that allow users to parameterize the error bound and affect the degree of speedup.

- If the solution quality guarantee is not required, then one can also use Max-Sum, MGM, or DSA. Their performance depends on a number of factors: If the problem has large domain sizes, MGM and DSA often outperform Max-Sum, since the memory and computational complexities of Max-Sum grows exponentially with the domain size. However, if the problem has small induced widths (for instance, when its constraint graph is acyclic), then Max-Sum is very efficient. It is even guaranteed to find optimal solutions when the induced width is 1. In general, Max-Sum tends to find solutions of good quality especially when considering its recent improvements (e.g., (Zivan, Parash, Cohen, Peled, & Okamoto, 2017)).

- If the problem has hard constraints (i.e., certain value combinations are prohibited), then the sampling algorithms (i.e., DUCT and D-Gibbs) are not recommended as they are not able to handle such problems. They require the cost functions to be smooth, and exploit that characteristic to explore the search space. Thus, one is restricted to search- or inference-based algorithms.

- In general, MGM and DSA are good robust benchmarks as they tend to find reasonably high quality solutions in practice. However, if specific problem characteristics are known, such as the ones discussed above, then certain algorithms may be able to exploit them to find better solutions.

4.5 Notable Variant: Asymmetric DCOPs

Asymmetric DCOPs (Grinshpoun, Grubshtein, Zivan, Netzer, & Meisels, 2013) are used to model multi-agent problems where agents controlling variables in the scope of a cost function can incur to different costs, given a fixed join assignment. Such a problem cannot be naturally represented by classical DCOPs, which require that all agents controlling variables participating in a cost function incur to the same cost as each other.

4.5.1 DEFINITION

An *Asymmetric DCOP* is a tuple $\langle \mathbf{A}, \mathbf{X}, \mathbf{D}, \mathbf{F}, \alpha \rangle$, where $\mathbf{A}, \mathbf{X}, \mathbf{D}$, and α are as defined in Definition 4.1, and each cost function $f_i \in \mathbf{F}$ is defined as: $f_i : \mathsf{X}_{x_j \in \mathbf{x}^i} D_j \times \alpha(f_i) \to (\mathbb{R}^+ \cup \{\bot\})$.

In other words, an Asymmetric DCOP is a DCOP where the cost that an agent incurs from a cost function may differ from the cost that another agent incurs from the same cost function.

As costs for participating agents may differ from each other, the goal in Asymmetric DCOPs is different from the goal in classical DCOPs. Given a cost function $f_j \in \mathbf{F}$ and complete assignment σ, let $f_j(\sigma, a_i)$ denote the cost incurred by agent a_i from cost function f_j with complete assignment σ. Then, the goal in Asymmetric DCOPs is to find the solution σ^*:

$$\sigma^* := \underset{\sigma \in \Sigma}{\operatorname{argmin}} \sum_{f_j \in \mathbf{F}} \sum_{a_i \in \alpha(f_j)} f_j(\sigma_{\mathbf{x}^j}, a_i) \tag{2}$$

As in classical DCOPs, solving Asymmetric DCOPs is NP-hard. In particular, it is possible to reduce any Asymmetric DCOP to an equivalent classical DCOP by introducing a polynomial number of variables and constraints, as described in the next section.

4.5.2 RELATION TO CLASSICAL DCOPs

One way to solve MAS problems with asymmetric costs via classical DCOPs is through the *Private Event As Variables (PEAV)* model (Maheswaran et al., 2004a). It can capture asymmetric costs by introducing, for each agent, as many "mirror" variables as the number of variables held by neighboring agents. The consistency with the neighbors' state variables is imposed by a set of equality constraints. However, this formalism suffers from scalability problems, as it may result in a significant increase in the number of variables in a DCOP. In addition, Grinshpoun et al. (2013) showed that most of the existing incomplete classical DCOP algorithms cannot be used to effectively solve Asymmetric DCOPs, even when the problems are reformulated through the PEAV model. They show that such algorithms are unable to distinguish between different solutions that satisfy all hard constraints, resulting in a convergence to one of those solutions and the inability to escape that local optimum. Therefore, it is important to design specialized algorithms to solve Asymmetric DCOPs.

4.5.3 ALGORITHMS

The current research direction in the design of Asymmetric DCOP algorithms has focused on adapting existing classical DCOP algorithms to handle the asymmetric costs. Asymmetric DCOPs require that each agent, whose variables participate in a cost function, coordinate the aggregation of their individual costs. To do so, two approaches have been identified (Brito, Meisels, Meseguer, & Zivan, 2009):

- A *two-phase strategy*, where only one side of the constraint (i.e., the cost induced by one agent) is considered in the first phase. The other side(s) (i.e., the cost induced by the other agent(s)) is considered in the second phase once a complete assignment is produced. As a result, the costs of all agents are aggregated.
- A *single-phase strategy*, which requires a systematic check of each side of the constraint before reaching a complete assignment. Checking each side of the constraint is often referred to as *back checking*, a process that can be performed either synchronously or asynchronously.

COMPLETE ALGORITHMS

SyncABB-2ph (Grinshpoun et al., 2013). *Synchronous Asymmetric Branch and Bound - 2-phase (SyncABB-2ph)* is a complete, synchronous, search-based algorithm that extends SyncBB with the

two-phase strategy. Phase 1 emulates SyncBB, where each agent considers the values of its cost functions with higher-priority agents. Phase 2 starts once a complete assignment is found. During this phase, each agent aggregates the sides of the cost functions that were not considered during Phase 1 and verifies that the known bound is not exceeded. If the bound is exceeded, Phase 2 ends and the agents restart Phase 1 by backtracking and resuming the search from the lower priority agent that exceeded the bound. The worst case runtime, memory, and communication requirements of this algorithm are the same as those of SyncBB.

SyncABB-1ph (Grinshpoun et al., 2013; Levit, Grinshpoun, Meisels, & Bazzan, 2013). *Synchronous Asymmetric Branch and Bound - 1-phase (SyncABB-1ph)* is a complete, synchronous, search-based algorithm that extends SyncBB with the one-phase strategy. Each agent, after having extended the CPA, updates the bound with its local cost associated to the cost functions involving its variables – as done in SyncBB. In addition, the CPA is sent back to the assigned agents to update its bound via a sequence of back checking operations. The worst case runtime, memory, and communication requirements of this algorithm are the same as those of SyncBB.

ATWB (Grinshpoun et al., 2013). The *Asymmetric Two-Way Bounding (ATWB)* algorithm is a complete, asynchronous, search-based algorithm that extends AFB to accommodate both forward bounding and backward bounding. The forward bounding is performed analogously to AFB. The backward bounding, instead, is achieved by sending copies of the CPA backward to the agents whose assignments are included in the CPA. Similar to what is done in AFB, agents that receive a copy of the CPA compute their estimates and send them forward to the assigning agent. The worst case runtime, memory, and communication requirements of this algorithm are the same as those of AFB.

INCOMPLETE ALGORITHMS

ACLS (Grinshpoun et al., 2013). *Asymmetric Coordinated Local Search (ACLS)* is an incomplete, synchronous, search-based algorithm that extends DSA. After a random value initialization, each agent exchanges its values with all its neighboring agents. At the end of this step, each agent identifies all possible improving assignments for its own variables, given the current neighbors choices. Each agent then selects one such assignments, according to the distribution of gains (i.e., reductions in costs) from each proposal assignment, and exchanges it with its neighbors. When an agent receives a proposal assignment, it responds with the evaluation of its side of the cost functions, resulting from its current assignment and the proposal assignments of the other agents participating in the cost function. After receiving the evaluations from each of its neighbors, each agent estimates the potential gain or loss derived from its assignment, and commits to a change with a given probability, similar to agents in DSA, to escape from local minima. The worst case runtime, memory, and communication requirements of this algorithm are the same as those of DSA.

MCS-MGM (Grinshpoun et al., 2013). *Minimal Constraint Sharing MGM (MCS-MGM)* is an incomplete, synchronous, search-based algorithm that extends MGM by considering each side of the cost function. Like MGM, the agents operate in an iterative fashion, where they exchange their current values at the start of each iteration. Afterwards, each agent sends the cost for its side of each cost function to its neighboring agents that participate in the same cost function.[5] Upon re-

---

5. This is a version of the algorithm with a guarantee that it will converge to a local optima. In the original version of the algorithm, which does not have such guarantee, each agent sends the cost only if its gain with the neighbor's new values is larger than the neighbor's last known gain.

ceiving this information, each agent knows the total cost for each cost function – by adding together the value of both sides of the cost function. Therefore, like in MGM, the agent can calculate the maximum gain (i.e., maximum reduction in costs) if it changes its values, and will send this information to all its neighbors. Upon receiving the gains of its neighbors, each agent changes its value if its gain is the largest among its neighbors. The worst case runtime, memory, and communication requirements of this algorithm are the same as those of MGM.

4.6 Notable Variant: Multi-objective DCOPs

Multi-objective Optimization (MOO) (Miettinen, 1999; Marler & Arora, 2004) aims at solving problems involving more than one objective function to be optimized simultaneously. In a MOO problem, optimal decisions need to accommodate potentially conflicting objectives. *Multi-objective DCOPs* extend MOO problems and DCOPs (Delle Fave, Stranders, Rogers, & Jennings, 2011).

4.6.1 DEFINITION

A *Multi-objective DCOP (MO-DCOP)* is a tuple $\langle \mathbf{A}, \mathbf{X}, \mathbf{D}, \vec{\mathbf{F}}, \alpha \rangle$, where $\mathbf{A}, \mathbf{X}, \mathbf{D}$, and α are as defined in Definition 4.1, and $\vec{\mathbf{F}} = [F_1, \ldots, F_h]^T$ is a vector of multi-objective functions, where each F_i is a set of cost functions f_j as defined in Definition 4.1. For a complete assignment σ of a MO-DCOP, let the cost for σ according to the i^{th} multi-objective optimization function set F_i, where $1 \leq i \leq h$, be

$$F_i(\sigma) := \sum_{f_j \in F_i} f_j(\sigma_{\mathbf{x}^j}) \tag{3}$$

The goal of a MO-DCOP is to find a complete assignment σ^* such that:

$$\sigma^* := \operatorname*{argmin}_{\sigma \in \Sigma} \vec{\mathbf{F}}(\sigma) = \operatorname*{argmin}_{\sigma \in \Sigma} [F_1(\sigma), \ldots, F_h(\sigma)]^T \tag{4}$$

where $\vec{\mathbf{F}}(\sigma)$ is a *cost vector* for the MO-DCOP. A solution to a MO-DCOP involves the optimization of a set of partially-ordered assignments. The above definition considers point-wise comparison of vectors—i.e., $\vec{\mathbf{F}}(\sigma) \leq \vec{\mathbf{F}}(\sigma')$ if $F_i(\sigma) \leq F_i(\sigma')$ for all $1 \leq i \leq h$. Typically, there is no single global solution where all the objectives are optimized at the same time. Thus, solutions of a MO-DCOP are characterized by the concept of *Pareto optimality*, which can be defined through the concept of *dominance*:

Definition 1 (Dominance) *A solution $\sigma \in \Sigma$ is* dominated *by a solution $\sigma' \in \Sigma$ iff $\vec{\mathbf{F}}(\sigma') \leq \vec{\mathbf{F}}(\sigma)$ and $F_i(\sigma') < F_i(\sigma)$ for at least one F_i.*

Definition 2 (Pareto Optimality) *A solution $\sigma' \in \Sigma$ is* Pareto optimal iff *it is not dominated by any other solution.*

Therefore, a solution is Pareto optimal *iff* there is no other solution that improves at least one objective function without deteriorating the cost of another function. Another important concept is the *Pareto front*:

Definition 3 (Pareto Front) *The* Pareto front *is the set of all cost vectors of all Pareto optimal solutions.*

Solving a MO-DCOP is equivalent to finding the Pareto front. However, even for tree-structured MO-DCOPs, the size of the Pareto front may be exponential in the number of variables.[6] Thus, multi-objective algorithms often provide solutions that may not be Pareto optimal but may satisfy other criteria that are significant for practical applications. A widely-adopted criterion is that of *weak Pareto optimality*:

Definition 4 (Weak Pareto Optimality) *A solution* $\sigma' \in \Sigma$ *is* weakly Pareto optimal iff *there is no other solution* $\sigma \in \Sigma$ *such that* $\vec{\mathbf{F}}(\sigma) < \vec{\mathbf{F}}(\sigma')$.

In other words, a solution is weakly Pareto optimal if there is no other solution that improves *all* of the objective functions simultaneously. An alternative approach to Pareto optimality is one that uses the concept of *utopia points*:

Definition 5 (Utopia Point) *A cost vector* $\vec{\mathbf{F}}^\circ = [F_1^\circ, \ldots, F_h^\circ]^T$ *is a* utopia point iff $F_i^\circ = \min_{\sigma \in \Sigma} F_i(\sigma)$ *for all* $1 \leq i \leq h$.

In other words, a utopia point is the vector of costs obtained by independently optimizing h DCOPs, each associated to one objective of the multi-objective function vector. In general, $\vec{\mathbf{F}}^\circ$ is unattainable. Therefore, different approaches focus on finding a *compromise solution* (Salukvad, 1971), which is a Pareto optimal solution that is *close* to the utopia point. The concept of *closeness* is dependent on the approach adopted.

Similar to their centralized counterpart, MO-DCOPs have been shown to be NP-hard (their decision versions), and #P-hard (the related counting versions), and to have exponentially many non-dominated points (Glaßer, Reitwießner, Schmitz, & Witek, 2010).

4.6.2 ALGORITHMS

This section categorizes the proposed MO-DCOP algorithms into two classes: *complete* and *incomplete* algorithms, according to their ability to find the complete set of Pareto optimal solutions or only a subset of it.

COMPLETE ALGORITHMS

MO-SBB (Medi, Okimoto, & Inoue, 2014). *Multi-objective Synchronous Branch and Bound (MO-SBB)* is a complete, synchronous, search-based algorithm that extends SyncBB. It uses an analogous search strategy to that of the mono-objective SyncBB: After establishing a complete ordering, MO-SBB agents extend a CPA with their own value assignments and the current associated cost vectors. Once a non-dominated solution is found, it is broadcasted to all agents, which add the solution to a list of global bounds. Thus, agents maintains an approximation of the Pareto front, which is used to bound the exploration, and extend the CPA only if the new partial assignment is not dominated by solutions in the list of global bounds. When the algorithm terminates, it returns the set of Pareto optimal solutions obtained by filtering the list of global bounds by dominance. The worst case runtime and communication requirements of this algorithm are the same as those of SyncBB. In terms of memory requirement, each MO-SBB agent needs $O(np)$ amount of memory, where p is the size of the Pareto set.

---

6. In the worst case, every possible solution is a Pareto optimal solution.

Pseudo-tree Based Algorithm (Matsui, Silaghi, Hirayama, Yokoo, & Matsuo, 2012). The proposed algorithm is a complete, asynchronous, search-based algorithm that extends ADOPT. It introduces the notion of boundaries on the vectors of multi-objective values, which extends the concept of lower and upper bounds to vectors of values. The proposed approach starts with the assumption that $|F_1| = \cdots = |F_h| = k$. Furthermore, the cost functions within each F_i are sorted according to a predefined ordering, and for each $1 \leq j \leq k$, the scope of f_j^i (i.e., the j^{th} function in F_i) is the same for each i (i.e., all functions in the same position in different F_i have the same scope). Thus, without loss of generality, the notation $\mathbf{x}^j$ will be used to refer to the scope of f_j^i.

Given a complete assignment σ, for $1 \leq j \leq k$, let $\vec{\sigma}_j = [\sigma_{f_j}^1, \ldots, \sigma_{f_j}^h] = [f_j^1(\sigma_{\mathbf{x}^j}), \ldots, f_j^h(\sigma_{\mathbf{x}^j})]$ be the vector of cost values. The notion of non-dominance is applied to these vectors, where a vector $\vec{\sigma}_j$ is non-dominated *iff* there is no other vector $\vec{\sigma}_j'$ such that $ub(\sigma_j'^s) \leq lb(\sigma_j^s)$ for all $1 \leq s \leq h$ and $ub(\sigma_j'^s) < lb(\sigma_j^s)$ for at least one s. The algorithm uses the notion of non-dominance for bounded vectors to retain exclusively non-dominated vectors.

The worst case runtime and communication requirements of this algorithm are the same as those of ADOPT. In terms of memory requirement, each agent needs $O(np)$ amount of memory. However, notice that the number of combinations of cost vectors grows exponentially with the number of tuples of cost values, in the worst case. This algorithm has also been extended to solve Asymmetric MO-DCOPs (Matsui, Silaghi, Hirayama, Yokoo, & Matsuo, 2014), which is an extension of both Asymmetric DCOPs and MO-DCOPs.

INCOMPLETE ALGORITHMS

B-MOMS (Delle Fave et al., 2011). *Bounded Multi-objective Max-Sum (B-MOMS)* is an incomplete, asynchronous, inference-based algorithm, and was the first MO-DCOP algorithm introduced. It extends Bounded Max-Sum to compute bound approximations for MO-DCOPs. It consists of three phases. The *Bounding Phase* generates an acyclic subgraph of the multi-objective factor graph, using a generalization of the maximum spanning tree problem to vector weights. During the *Max-sum Phase*, the agents coordinate to find the Pareto optimal set of solutions to the acyclic factor graph generated in the bounding phase. This is achieved by extending the addition and marginal maximization operators adopted in Max-Sum to the case of multiple objectives. Finally, the *Value Propagation Phase* allows agents to select a consistent variable assignment, as there may multiple Pareto optimal solutions. The bounds provided by the algorithm are computed using the notion of utopia points.

The worst case runtime requirement of this algorithm is the same as those of Max-Sum. In terms of communication requirement, the number of messages sent is also like Max-Sum, but the size of each message is now $O(pd^n)$. In terms of memory requirement, each B-MOMS agent needs $O(pd^n)$ amount of memory to store and process the messages received.

DP-AOF (Okimoto, Clement, & Inoue, 2013). *Dynamic Programming based on Aggregate Objective Functions (DP-AOF)* is an incomplete, synchronous, inference-based algorithm. It adapts the AOF technique (Miettinen, 1999), designed to solve centralized multi-objective optimization problems, to solve MO-DCOPs. Centralized AOF adopts a scalarization to convert a MOO problem into a single objective optimization. This is done by assigning weights $(\alpha_1, \ldots, \alpha_h)$ to each of the cost functions in the objective vector $[F_1, \ldots, F_h]^T$ such that $\sum_{i=1}^h \alpha_i = 1$ and $\alpha_i > 0$ for all $1 \leq i \leq h$. The resulting mono-objective function $\sum_{i=1}^h \alpha_i F_i$ can be solved using any mono-objective optimization technique with guarantee to find a Pareto optimal solution (Miettinen, 1999).

DP-AOF proceeds in two phases. First, it computes the utopia point $\vec{\mathbf{F}}^\circ$ by solving as many mono-objective DCOPs as the number of objective functions in the MO-DCOP. DP-AOF uses DPOP to solve these mono-objective DCOPs. It then constructs a new problem building upon the solutions obtained from the first phase. Such a problem is used to assign weights to each objective function of the MO-DCOP to construct the new mono-objective function in the same way as centralized AOF, which then can be solved optimally. The worst case runtime, memory, and communication requirements of this algorithm are the same as those of DPOP, except that the number of operations and the number of messages are larger by a factor of h since it runs DPOP h times to solve the h mono-objective DCOPs.

MO-DPOP$_{L_p}$ (Okimoto, Schwind, Clement, & Inoue, 2014). *Multi-objective L_p-norm based Distributed Pseudo-tree Optimization Procedure (MO-DPOP$_{Lp}$)* is an incomplete, synchronous, inference-based algorithm. It adapts DPOP using a scalarization measure based on the L_p-norm to find a subset of the Pareto front of a MO-DCOP. Similar to DP-AOF, the algorithm proceeds in two phases. Its first phase is the same as the first phase of DP-AOF: It solves h mono-objective DCOPs using DPOP to find the utopia point $\vec{\mathbf{F}}^\circ$. In the second phase, the agents coordinate to find a solution that minimizes the distance from $\vec{\mathbf{F}}^\circ$ according to the L_p-norm. The algorithm is guaranteed to find a Pareto optimal solution only when the L_1-norm (Manhattan norm) is adopted. In this case, MO-DPOP$_{L_1}$ finds a Pareto optimal solution that minimizes the average cost values of all objectives. The worst case runtime, memory, and communication requirements of this algorithm are the same as those of DP-AOF.

DIPLS (Wack, Okimoto, Clement, & Inoue, 2014). *Distributed Iterated Pareto Local Search (DIPLS)* is an incomplete, synchronous, search-based algorithm. It extends the *Pareto Local Search (PLS)* algorithm (Paquete, Chiarandini, & Sttzle, 2004), which is a hill climbing algorithm designed to solve centralized multi-objective optimization problems, to solve MO-DCOPs. The idea behind DIPLS is to evolve an initial solution toward the Pareto front. To do so, it starts from an initial set of random assignments, and applies PLS iteratively to generate new non-dominated solutions. DIPLS requires a total ordering of agents and elects one agent as the *controller*. At each iteration, the controller filters the set of solutions by dominance and broadcasts them to the agents in the MO-DCOP. Upon receiving a solution, an agent generates a list of *neighboring solutions* by modifying the assignments of the variables that it controls, and sends them back to the controller. When the controller receives the messages from all agents, it proceeds to filter (by dominance) the set of solutions received, and if a new non-dominated solution is found, it repeats the process.

The worst case runtime of this algorithm is $O(\ell k p)$ as the controller agent is required to check the dominance of the newly generated solutions at each iteration. In terms of memory requirement, DIPLS agents use $O(np)$ space to store the Pareto front. Finally, in terms of communication requirement, the controller agent broadcasts messages that contain the current Pareto front. Thus, the message size is $O(np)$.

5. Dynamic DCOPs

Within a real-world MAS application, agents often act in dynamic environments that evolve over time. For instance, in a disaster management search and rescue scenario, new information (e.g., the number of victims in particular locations or priorities on the buildings to evacuate) typically becomes available in an incremental manner. Thus, the information flow modifies the environment

over time. To cope with such a requirement, researchers have introduced the *Dynamic DCOP (D-DCOP)* model, where cost functions can change during the problem solving process, agents may fail, and new agents may be added to the DCOP being solved. With respect to the categorization described in Section 3, in the D-DCOP model, the agents are fully cooperative and they have deterministic behavior and total knowledge. On the other hand, the environment is dynamic and deterministic.

5.1 Definition

The *Dynamic DCOP (D-DCOP)* model is defined as a sequence of classical DCOPs: $\mathcal{D}_1, \ldots, \mathcal{D}_T$, where each $\mathcal{D}_t = \langle \mathbf{A}^t, \mathbf{X}^t, \mathbf{D}^t, \mathbf{F}^t, \alpha^t \rangle$ is a DCOP representing the problem at time step t, for $1 \leq t \leq T$. The goal in a D-DCOP is to solve the DCOP at each time step optimally. By assumption, the agents have total knowledge about their current environment (i.e., the current DCOP) but they are unaware of changes to the problem in future time steps.

In a dynamic system, agents are required to adapt as fast as possible to environmental changes. *Stability* (Dijkstra, 1974; Verfaillie & Jussien, 2005) is a core algorithmic concept in which an algorithm seeks to minimize the number of steps that it requires to converge to a solution each time the problem changes. In such a context, these converged solutions are also called *stable solutions*. *Self-stabilization* is a related concept derived from the area of fault-tolerance:

Definition 6 (Self-stabilization) *A system is* self-stabilizing iff *the following two properties hold:*

(i) Convergence: *The system reaches a stable solution in a finite number of steps, starting from any given state. In the DCOP context, this property expresses the ability of the agents to coordinate a joint assignment for their variables that optimizes the problem at time step $t+1$, starting from an assignment of the problem's variables at time step t.*

(ii) Closure: *The system remains in a stable solution, provided that no changes in the environment happens. In the DCOP context, this means that agents do not change the assignment for their variables after converging to a solution.*

Solving D-DCOPs is NP-hard, as it requires to solve each DCOP of the D-DCOP independently.

5.2 Algorithms

In principle, one could use classical DCOP algorithms to solve the DCOP $\mathcal{D}_t$ at each time step $1 \leq t \leq T$. However, the dynamic environment evolution encourages firm requirements on the algorithm design in order for the agents to respond automatically and efficiently to environmental changes over time. In particular, D-DCOP algorithms often follow the *self-stabilizing* property. As in the previous sections, the algorithms are categorized as being either complete or incomplete, according to their ability to determine the optimal solution at each time step.

5.2.1 COMPLETE ALGORITHMS

S-DPOP (Petcu & Faltings, 2005c). *Self-stabilizing DPOP (S-DPOP)* is a synchronous, inference-based algorithm that extends DPOP to handle dynamic environments. It is composed of three self-stabilizing phases: *(i)* A self-stabilizing DFS pseudo-tree generation, whose goal is to create and maintain a DFS pseudo-tree structure; *(ii)* A self-stabilizing algorithm for the UTIL propagation phase; and *(iii)* A self-stabilizing algorithm for the VALUE propagation phase. These procedures

work as in DPOP and they are invoked whenever any change in the DCOP problem sequence is revealed. Additionally, Petcu and Faltings (2005c) discuss self-stabilizing extensions that can be used to provide guarantees about the way the system transitions from a valid state to the next, after an environment change.

The worst case runtime, memory, and communication requirements of this algorithm to solve the DCOP at each time step are the same as those of DPOP. Additionally, upon changes to the problem, S-DPOP stabilizes after at most τ UTIL messages and k VALUE messages, where τ is the depth of the pseudo-tree and k is the number of cost functions of the problem.

I-ADOPT and I-BnB-ADOPT (Yeoh, Varakantham, Sun, & Koenig, 2011). *Incremental Any-space ADOPT (I-ADOPT)* and *Incremental Any-space BnB-ADOPT (I-BnB-ADOPT)* are asynchronous, search-based algorithms that extend ADOPT and BnB-ADOPT, respectively. In the incremental any-space versions of the algorithms, each agent maintains bounds for multiple contexts; in contrast, agents in ADOPT and BnB-ADOPT maintain bounds for one context only. By doing so, when solving the next DCOP in the sequence, agents may reuse the bounds information computed in the previous DCOP. In particular, the algorithms identify *affected agents*, which are agents that cannot reuse the information computed in the previous iterations, and they recompute bounds exclusively for such agents.

The worst case runtime and communication requirements of this algorithm to solve the DCOP at each time step are the same as those of ADOPT. However, since these algorithms have the any-space property, their minimal memory requirements are the same as those of ADOPT but they can use more memory, if available, to speed up the algorithms.

5.2.2 INCOMPLETE ALGORITHMS

SBDO (Billiau, Chang, & Ghose, 2012a). *Support Based Distributed Optimization (SBDO)* is an asynchronous search-based algorithm that extends the *Support Based Distributed Search* algorithm (Harvey, Chang, & Ghose, 2007) to the multi-agent case. It uses two types of messages: *is-good* and *no-good*. Is-good messages contain an ordered partial assignment and are exchanged among neighboring agents upon a change in their value assignments. Each agent, upon receiving a message, decides what value to assign to its own variables, attempting to minimize their local costs, and communicates such decisions to its neighboring agents via is-good messages. No-good messages are used in response to violations of hard constraints, or in response to obsolete assignments. A no-good message is augmented with a *justification*, that is, the set of hard constraints that are violated, and are saved locally within each agent. This information is used to discard partial assignments that are supersets of one of the known no-goods. The changes of the dynamic environment are communicated via messages, which are sent from the *environment* to the agents. In particular, changes in hard constraints require the update of all the justifications in all no-goods.

The worst case runtime, memory, and communication requirements of this algorithm are the same as those of SyncBB each time the problem changes.

FMS (Ramchurn, Farinelli, Macarthur, & Jennings, 2010). *Fast Max-Sum (FMS)* is an asynchronous inference-based algorithm that extends Max-Sum to the Dynamic DCOP model. As in Max-Sum, the algorithm operates on a factor graph. Solution stability is maintained by recomputing only those factors that changed between the previous DCOP $\mathcal{D}_{t-1}$ and the current DCOP $\mathcal{D}_t$. Ramchurn et al. (2010) exploit domain-specific properties in a task allocation problem to reduce the

number of states over which each factor has to compute its solution. In addition, FMS is able to efficiently manage addition or removal of tasks (e.g., factors), by performing message propagation exclusively on the factor graph regions that are affected by such topological changes. The worst case runtime, memory, and communication requirements of this algorithm to solve the DCOP at each time step are the same as those of Max-Sum.

FMS has been extended in several ways. *Bounded Fast Max-Sum* provides bounds on the solution found, as well as it guarantees self-stabilization (Macarthur, Farinelli, Ramchurn, & Jennings, 2010). *Branch-and-Bound Fast Max-Sum (BnB-FMS)* extends FMS providing online domain pruning using a branch-and-bound technique (Macarthur, Farinelli, Ramchurn, & Jennings, 2011).

5.3 Notable Variants: D-DCOPs with Commitment Deadlines or Markovian Properties

We now describe several notable variants of D-DCOPs and their corresponding algorithms.

RS-DPOP (Petcu & Faltings, 2007b). In this proposed model, agents have *commitment deadlines* and *stability constraints*. In other words, some of the variables may be unassigned at a given point in time, while others must be assigned within a specific deadline. Commitment deadlines are either hard or soft. *Hard commitments* model irreversible processes. When a hard committed variable is assigned, its value cannot be changed. *Soft commitments* model contracts with penalties. If a soft committed variable x_i^t has been assigned at time step t, its value can be changed at time step $t' > t$, at the price of a cost penalty. These costs are modeled via *stability constraints*, which are defined as binary relations $s_i : D_i \times D_i \to \mathbb{R}^+$, representing the cost of changing the value of variable x_i from time step t to time step $t + 1$. Given the set of stability constraints $\mathbf{S} \subseteq \mathbf{F}$, at each time step t, the goal is to find a solution σ_t^*:

$$\sigma_t^* := \operatorname*{argmin}_{\sigma \in \Sigma} \left(\mathbf{F}_g(\sigma) + \sum_{s_j \in \mathbf{S}} s_j(\sigma_{t-1}^*(x_j), \sigma(x_j)) \right).$$

The latter term accounts for the penalties associated to the value assignment updates for the soft committed variables. RS-DPOP has the same order complexity as S-DPOP.

To solve this problem, Petcu and Faltings (2007b) extended S-DPOP to RS-DPOP.[7] Like S-DPOP, it is a synchronous, inference-based algorithm. Unlike S-DPOP, it's UTIL and VALUE propagation phases now take into account the commitment deadlines. The worst case runtime, memory, and communication requirements of this algorithm to solve the DCOP at each time step are the same as those of S-DPOP.

Distributed Q-learning and R-learning (Nguyen, Yeoh, Lau, Zilberstein, & Zhang, 2014). In this proposed model, called *Markovian Dynamic DCOPs (MD-DCOPs)*, the DCOP in the next time step $\mathcal{D}_{t+1}$ depends on the solution (i.e., assignment of all variables) adopted by the agents for the DCOP in the current time step $\mathcal{D}_t$. However, the transition function between these two DCOPs are not known to the agents and the agents must, thus, learn them. The *Distributed Q-learning and R-learning* algorithms are synchronous reinforcement-learning-based algorithms that extend the centralized Q-learning (Abounadi, Bertsekas, & Borkar, 2001) and centralized R-learning (Schwartz, 1993; Mahadevan, 1996) algorithms. Each agent maintains Q-values and R-values for each σ_{t-1}, d_i^t pair, where σ_{t-1} is the solution for the DCOP $\mathcal{D}_{t-1}$ and d_i^t is the value of its variables

---

7. The full name of the algorithm was not provided by Petcu and Faltings (2007b).

in the cost function $f_i^t \in \mathbf{F^t}$. These Q- and R-values represent the predicted cost the agent will incur if it assigns its variables values according to d_i^t when σ_{t-1} is the previous solution. The agents repeatedly refine these values at every time step and choose the values with the minimum Q- or R-value at each time step.

The worst case runtime, communication, and memory requirements of these two algorithms to solve the DCOP at each time step are the same as those of DPOP, as they use DPOP as a subroutine to update the Q- and R-values. The exception is that agents in the Distributed Q-learning algorithm also broadcast their value assignments at each time step to all other agents. Thus, they send $O(m^2)$ messages in each time step instead of the $O(m)$ complexity of DPOP.[8]

A related model is the *Proactive Dynamic DCOPs (PD-DCOPs)* (Hoang, Fioretto, Hou, Yokoo, Yeoh, & Zivan, 2016; Hoang, Hou, Fioretto, Yeoh, Zivan, & Yokoo, 2017), where the transition functions between two subsequent DCOPs are known and can be exploited by the resolution process. Additionally, another key difference between these two models is that the DCOP in the next time step $\mathcal{D}_{t+1}$ does not depend on the solution in the current time step, but instead depends on the values of the random variables at the current time step. Researchers have introduced a number of offline proactive and online reactive algorithms to solve this problem (Hoang et al., 2016, 2017).

6. Probabilistic DCOPs

The DCOP models discussed so far can model MAS problems in deterministic environments. However, many real-world applications are characterized by environments with a stochastic behavior. In other words, there are exogenous events that can influence the outcome of an agent's action. For example, the weather conditions or the state of a malfunctioning device can affect the cost of an agent's action. To cope with such scenarios, researchers have introduced *Probabilistic DCOP (P-DCOP)* models, where the uncertainty in the state of the environment is modeled through stochasticity in the cost functions. With respect to the DCOP categorization described in Section 3, in the P-DCOP model, the agents are fully cooperative and have a deterministic behavior. Additionally, the environment is static and stochastic. While a large body of research has focused on problems where agents have *total knowledge*, this section includes a discussion of a subclass of P-DCOPs where the agents' knowledge of the environment is limited, and the agents must balance the *exploration* of the unknown environment and the *exploitation* of the known costs.

6.1 Definition

A common strategy to model uncertainty is to augment the outcome of the cost functions with a stochastic character (Atlas & Decker, 2010; Stranders, Delle Fave, Rogers, & Jennings, 2011; Nguyen, Yeoh, & Lau, 2012). Another method is to introduce additional random variables as input to the cost functions, which simulate exogenous uncontrollable traits of the environment (Léauté & Faltings, 2009, 2011; Wang, Sycara, & Scerri, 2011). To cope with such a variety, this section introduces the *Probabilistic DCOP (P-DCOP)* model, which generalizes the proposed models of uncertainty. A P-DCOP is defined by a tuple $\langle \mathbf{A}, \mathbf{X}, \mathbf{D}, \mathbf{F}, \alpha, \mathcal{I}, \Omega, \mathbf{P}, \mathcal{E}, \mathcal{U} \rangle$, where $\mathbf{A}$ and $\mathbf{D}$ are as defined in Definition 4.1. In addition,

- $\mathbf{X}$ is a mixed set of decision variables and random variables.

---

8. A single broadcast message is counted as m peer-to-peer messages, where m is the number of agents in the problem.

- $\mathcal{I} = \{r_1, \ldots, r_q\} \subseteq \mathbf{X}$ is a set of *random variables* modeling uncontrollable stochastic events, such as weather or a malfunctioning device.
- $\mathbf{F} = \{f_1, \ldots, f_k\}$ is the set of cost functions, each defined over a mixed set of decision variables and random variables, and such that each value combination of the decision variables on the cost function results in a probability distribution. As a result, f_i is itself a random variable, given the local value assignment $\sigma_{\mathbf{x}^i \setminus \mathcal{I}}$ and a realization for the random variables involved in f_i.
- $\alpha : \mathbf{X} \setminus \mathcal{I} \to \mathbf{A}$ is a mapping from decision variables to agents. Notice that random variables are not controlled by any agent, as their outcomes do not depend on the agents' actions.
- $\Omega = \{\Omega_1, \ldots, \Omega_q\}$ is the (possibly discrete) set of events for the random variables (e.g., the different weather conditions or stress levels a device is subjected to) such that each random variable $r_i \in \mathcal{I}$ takes values in Ω_i. In other words, Ω_i is the domain of random variable r_i.
- $\mathbf{P} = \{p_1, \ldots, p_q\}$ is a set of probability distributions for the random variables, such that $p_i : \Omega_i \to [0,1] \subseteq \mathbb{R}$ assigns a probability value to an event for r_i and $\int_{\omega \in \Omega_i} p_i(\omega) \, d\omega = 1$ for each random variable $r_i \in \mathcal{I}$.
- $\mathcal{E}$ is an *evaluator function* from random variables to real values, that, given an assignment of values to the decision variables, summarizes the distribution of the aggregated cost functions.
- $\mathcal{U}$ is a *utility function* that given a random variable returns an ordered set of different outcomes, and it is based on the decision maker preferences. This function is needed when the cost functions have uncertain outcomes and, thus, these distributions are not readily comparable.

The goal in a P-DCOP is to find a solution σ^*, that is, an assignment of values to all the decision variables, such that:

$$\sigma^* := \arg \min/\max_{\sigma \in \Sigma} \mathcal{E} \left[\sum_{f_i \in \mathbf{F}} \mathcal{U}\big(f_i(\sigma_{\mathbf{x}^i \setminus \mathcal{I}})\big) \right] \tag{5}$$

where argmin or argmax are selected depending on the algorithm adopted, $\sum$ is the operator that is used to aggregate the values from the functions $f_i \in \mathbf{F}$. Typically such an operator is a summation, however, to handle continuous distributions, other operators have been proposed.

The probability distribution over the domain of random variables $r_i \in \mathcal{I}$ is called a *belief*. An assignments of all random variables in $\mathcal{I}$ describes a (possible) *scenario* governed by the environment. As the random variables are not under the control of the agents, they *act* independently of the decision variables. Specifically, their beliefs are drawn from probability distributions. Furthermore, they are assumed to be independent of each other and, thus, they model independent sources of exogenous uncertainty.

The utility function $\mathcal{U}$ enables us to compare the uncertain cost outcomes of the cost functions. In general, the utility function is non-decreasing, that is, the lower the cost, the higher the utility. However, the utility function should be defined for the specific application of interest. For example, in farming, the utility increases with the amount of produce harvested. However, farmers may prefer a smaller but highly certain amount of produce harvested over a larger but highly uncertain and, thus, risky outcome.

The evaluation function $\mathcal{E}$ is used to summarize in one criterion the costs of a given assignment that depends on the random variables. A possible evaluation function is the *expectation* function: $\mathcal{E}[\cdot] = \mathbb{E}[\cdot]$.

Let us now introduce some concepts that are commonly adopted in the study of P-DCOPs.

Definition 7 (Convolution) *The* convolution *of the* probability density function (PDF) $f(x)$ *and* $g(x)$ *of two independent random variables X and Y is the integral of the product of the two functions after one is reversed and shifted:*

$$h(z) = (f * g)(z) := \int_{-\infty}^{\infty} f(\tau)\, g(z - \tau)\, d\tau = \int_{-\infty}^{\infty} f(z - \tau)\, g(\tau)\, d\tau \qquad (6)$$

It produces a new PDF $h(z)$ that defines the overlapping area between $f(x)$ and $g(y)$ as a function of the quantity that one of the original functions is translated by. In other words, the convolution is a method of determination of the sum of two random variables. The counterpart for the distribution of the sum $Z = X + Y$ of two independent discrete variables is:

$$P(Z = z) = \sum_{k=-\infty}^{\infty} P(X = k)P(Y = z - k). \qquad (7)$$

In a P-DCOP, the value returned by a function f_i, for an assignment on its scope $\mathbf{x}^i$, is a random variable V_i ($V_i \sim f_i(\mathbf{x}^i)$). Thus, the global value $\sum_{f_i \in \mathbf{F}} V_i$ is also a random variable, whose probability density function is the convolution of the PDFs of the individual V_i's. Thus, the concept of convolution of two PDFs in a P-DCOP is related to the summation of the utilities of two cost functions in classical DCOPs.

A common concept in optimization with uncertainty is that of *ranking* a set of random variables $\{r_1, r_2, \dots\}$ with *Cumulative PDFs (CDFs)* $\{F_1(x), F_2(x), \dots\}$. These distributions are also commonly called *lotteries*, a concept related to that of *stochastic dominance*, which is a form of stochastic ordering based on preference regarding outcomes. It refers to situations where a probability distribution over possible outcomes can be ranked as superior to another.

The first-order stochastic dominance refers to the situation when one lottery is unambiguously better than another:

Definition 8 (First-Order Stochastic Dominance) *Given two random variables r_i and r_j with CDFs $F_i(x)$ and $F_j(x)$, respectively, F_i first-order stochastically dominates F_j iff:*

$$F_i(x) \le F_j(x), \qquad (8)$$

for all x with a strict inequality over some interval.

If F_i first-order stochastically dominates F_j, then F_i necessarily has a strictly smaller expected value: $\mathbb{E}[F_i(x)] < \mathbb{E}[F_j(x)]$. In other words, if F_i dominates F_j, then the decision maker prefers F_i over F_j regardless of his utility function $\mathcal{U}$ is, as long as it is weakly increasing.

It is not always the case that one CDF will first-order stochastically dominate another. In such a case, one can use the second-order stochastic dominance to compare them. The latter refers to the situation when one lottery is unambiguously less *risky* than another:

Definition 9 (Second-Order Stochastic Dominance) *Given two random variables r_i and r_j with CDFs $F_i(x)$ and $F_j(x)$, respectively, F_i second-order stochastically dominates F_j iff:*

$$\int_{-\infty}^{c} F_i(x)\, dx \le \int_{-\infty}^{c} F_j(x)\, dx, \qquad (9)$$

for all c with a strict inequality for some values of c.

If F_i second-order stochastically dominates F_j, then $\mathbb{E}[F_i(x)] \leq \mathbb{E}[F_j(x)]$. If Equation 9 holds for all $c \geq c'$, for some sufficiently large c', then $\mathbb{E}[F_i(x)] = \mathbb{E}[F_j(x)]$. In this case, as both lotteries are equal in expectation, the decision maker prefers the lottery F_i, which has less variance and is, thus, less risky.

Another common concept in P-DCOPs is that of *regret*. In decision theory, regret expresses the negative *emotion* arising from learning that a different solution than the one adopted, would have had a more favorable outcome. In P-DCOPs the regret of a given solution is typically defined as the difference between its associated cost and that of the theoretical optimal solution. The notion of regret is especially useful in allowing agents to make robust decisions in settings where they have limited information about the cost functions.

An important type of regret is the *minimax regret*. Minimax regret is a decision rule used to minimize the possible loss for a worst case (i.e, maximum) regret. As opposed to the (expected) regret, minimax regret is independent of the probabilities of the various outcomes. Thus, minimax regret could be used when the probabilities of the outcomes are unknown or difficult to estimate.

Solving P-DCOPs is PSPACE-hard, as in general, the process is required to remember a solution for each possible state associated to the uncertain random variables. The study of complexity classes for P-DCOPs is largely unexplored. Thus, we foresee this as a potential direction for future research, in which particular focus could be given in determining fragments of P-DCOPs characterized by lower complexity than the one above.

6.2 Algorithms

Unlike for Classical DCOPs and Dynamic DCOPs, where the algorithms solve the same problem, P-DCOP algorithms approach the problem uncertainty in different ways and, thus, solve different variants of the problem. This is due to the greater modeling flexibility offered by the P-DCOP framework. As such, the proposed algorithms are often not directly comparable to one another. We categorize P-DCOP algorithms into *complete* and *incomplete* algorithms, according to their ability to guarantee to find the optimal solutions or not, for a given evaluator and utility functions. Unless otherwise specified the ordering operator in Equation 5 refers to the `argmax` operator.

6.2.1 COMPLETE ALGORITHMS

$\mathbb{E}$**[DPOP]** (Léauté & Faltings, 2011). $\mathbb{E}$[DPOP] is a synchronous, sampling-based and inference-based algorithm. It can be either complete or incomplete based on the $\mathbb{E}$[DPOP] variant used, and described below. $\mathbb{E}$[DPOP] uses a *collaborative sampling* strategy, where all agents concerned with a given random variable agree on a common sample set that will be used to estimate the PDF of that random variable. Agents performing collaborative sampling independently propose sample sets for the random variables influencing the variables they control, and elect one agent among themselves as responsible for combining the proposed sample sets into one. The algorithm is defined over P-DCOPs with $\mathcal{I} \neq \emptyset$ and deterministic cost function outcomes, that is, for each combination of values for the variables in $\mathbf{x}^i$, $f_i(\sigma_{\mathbf{x}^i \setminus \mathcal{I}})$ is a degenerate distribution (i.e., a distribution that results in a single value) and the utility function $\mathcal{U}$ is the identity function. $\mathcal{E}$ is an arbitrary evaluator function summing over all functions in $\mathbf{F}$.

$\mathbb{E}$[DPOP] builds on top of DPOP and proceeds in four phases: In Phase 1, the agents order themselves into a pseudo-tree ignoring the random variables. In Phase 2, the agents bind random variables to some decision variable. In Phases 3 and 4, the agents run the UTIL and VALUE

propagation phases like in DPOP except that random variables are sampled. Based on different strategies adopted in binding the random variables in Phase 2, the algorithm has two variants (Léauté & Faltings, 2009). In *Local-$\mathbb{E}$[DPOP]*, a random variable $r_i \in \mathcal{I}$ is assigned to each decision variable responsible for enforcing a constraint involving r_i. In this approach, the agents do not collaborate by exchanging information about how their utilities depend on the random variables. In contrast, *Global-$\mathbb{E}$[DPOP]* assigns r_i to the *lowest common ancestor* agent,[9] which is responsible for combining the proposed samples. While this additional information can produce higher-quality solutions, both algorithms are generally incomplete. One exception is when the evaluation function $\mathcal{E}$ adopted is *linear*, as in the case of the *expectation* function, in which case the algorithms are complete.

The worst case runtime, memory, and communication requirements of this algorithm are the same as those of DPOP. The exception is the message size of *Global-$\mathbb{E}$[DPOP]*, which is $O(d^{w^*}s^q)$, where s is the largest sample set size; the UTIL messages have this size when the root as well as all leaves of the pseudo-tree are constrained with all q random variables.

SD-DPOP (Nguyen et al., 2012). *Stochastic Dominance DPOP (SD-DPOP)* operates on a P-DCOP model where $\mathcal{I} = \emptyset$, $\mathcal{E}$ is the second order stochastic dominance criteria, $\mathcal{U}$ is the identity function, and Σ denotes the convolution of the distributions $f_i(\sigma_{\mathbf{x}^i \setminus \mathcal{I}})$. It is a complete synchronous inference-based algorithm that extends DPOP to solve P-DCOPs. Similar to DPOP, it has three phases. In Phase 1, like DPOP, it constructs a pseudo-tree. In Phase 2, instead of summing up costs, the agents convolve cost functions, and instead of propagating costs up the pseudo-tree, they propagate convolved cost functions. In Phase 3, like DPOP, the agents choose values for their variables. However, instead of choosing values that minimize the cost of their subtrees, the agents choose their values according to the second-order stochastic dominance criteria.

Like DPOP, SD-DPOP requires a linear number of messages. In addition, in SD-DPOP, VALUE messages contain each Pareto optimal value of the sending agent, and UTIL messages contain a representation of the cost function for each Pareto optimal solution and each combination of values of the parent and pseudo-parents of the sending agents. Thus, for continuous PDFs that could be represented by mean and variance, the message size is $O(pd^{w^*})$, where p is the size of the Pareto set. If the cost functions are represented by discretized bins, then the message size is $O(bpd^{w^*})$, where b is the maximum number of bins used to represent a cost function. The memory requirement of each SD-DPOP agent is $O(pd^{w^*})$ or $O(bpd^{w^*})$, similarly as above, as they need to store and process the messages received. The worst case runtime requirement and the number of messages sent by SD-DPOP are the same as those of DPOP.

6.2.2 INCOMPLETE ALGORITHMS

DNEA (Atlas & Decker, 2010). The P-DCOP model proposed by Atlas and Decker (2010) is characterized by uncertainty exclusively at the level of the outcome of the cost functions, and not due to random variables. Thus, $\mathcal{I} = \emptyset$. In addition, the utility function $\mathcal{U}$ is the identity function, while $\mathcal{E}$ is a given evaluator function (e.g., the expectation) for the functions $f_i \in \mathbf{F}$. In this settings, by employing the evaluation function, Atlas and Decker show that one can reduce the uncertainty associated to each cost function to the deterministic case. Thus, one can solve the proposed P-DCOP problems using classical DCOP approaches.

---

9. The agent that is separated by the smallest number of tree edges from all variables constrained with the given random variable.

In particular, they propose the *Distributed Neighbor Exchange Algorithm (DNEA)*, which is an incomplete, synchronous, search-based algorithm that is similar to DSA. Each agent starts by assigning a random value to each of its variables and sends this information to all its neighbors. Upon receiving the values of its neighbors, it computes a cost vector containing the costs for each possible combination of values for all its variables, under the assumption that its neighbors' values are those in the messages received. It then sends this cost vector to all its neighbors. Upon receiving the cost vector of its neighbors, it computes the best value for each of its variables, assigns those values to its variables probabilistically, and sends the assigned values to all its neighbors. This process repeats until a termination condition is satisfied.

The runtime requirement of this algorithm is $O(\ell(ld + d^2))$. In terms of communication requirement, the number of messages sent is $O(\ell n l)$, and the size of each message is $O(d)$. In terms of memory requirement, each DNEA agent needs $O(ld)$ amount of memory to store and process the messages received.

U-GDL (Stranders et al., 2011). The P-DCOP model proposed by Stranders et al. (2011) also assumes that the cost functions are not dependent on random variables. Thus, $\mathcal{I} = \emptyset$. Additionally, they assume that $\mathcal{E}$ is the expectation of the convolution (Σ) of the distributions $f_i(\sigma_{\mathbf{x}^i \setminus \mathcal{I}})$, and $\mathcal{U}$ is a given risk function. Stranders et al. propose the *Uncertain Generalized Distributive Law (U-GDL)* algorithm, which is an incomplete asynchronous inference-based algorithm similar to Max-Sum, and operates on acyclic graphs. A cyclic constraint graph G is converted into an acyclic graph $\hat{G}$ by merging variables until the new resulting graph contains no cycles. Merging two variables creates a new variable whose domain is the Cartesian product of the domain of the merged variables. U-DGL extends the *Generalized Distributive Law (GDL)* algorithm (Aji & McEliece, 2000) by redefining the $(\min, +)^{10}$ algebra to the setting where costs are random variables rather than scalars. The $+$ operator is extended to perform convolution of two random variables. To cope with the potential issue that not all PDFs are closed under convolution, Stranders et al. (2011) suggest to resort to sampling methods to approximate such operations. The $\min$ operator is defined to distribute over convolution and to select the minimal elements from a set of random variables based on their expected cost. In order to filter partial potential solutions that can never achieve global optimality, the authors introduce a first-order stochastic dominance condition, which is employed in the context of the $\min$ operator. They also discuss *necessary* and *sufficient* conditions for dominance, where the former discards all dominated solutions, but it might also discard some non-dominated solution – this is equivalent to using a classical DCOP algorithms to solve the P-DCOP model adopted in their work. The latter preserves optimal solutions, but retains, in general, sub-optimal ones as well.

The runtime requirement of this algorithm is $O(p\hat{d}^l)$, where $\hat{d}$ is the size of the largest domain of the merged variables in $\hat{G}$. In terms of memory requirement, each U-GDL agent needs $O(p\hat{d}^l)$ space to store all value combinations of neighboring variables for each solution in its current Pareto frontier. In terms of communication requirement, the number of messages sent is $O(\delta(\hat{G}))$, where $\delta(\hat{G})$ is the diameter of the resulting acyclic graph, and the size of each message is $O(p\hat{d})$ as agents need to send the current aggregated costs of all the agent's variable's values for each solution in its current Pareto frontier.

---

10. U-GDL was originally defined for maximization problems. Its presentation is adapted to minimization problems for consistency of the DCOP models objective presented in this survey.

6.3 Notable Variant: P-DCOPs with Partial Agent Knowledge

This section describes a class of Probabilistic DCOPs where agents have *partial* knowledge about the environment. In other words, the cost functions are only partially known and, therefore, agents may discover the unknown costs via *exploration* (Taylor, Jain, Tandon, Yokoo, & Tambe, 2011). The new model aims at capturing those domains where agents have an "explorative nature," i.e., one of the agents' goals is to acquire knowledge about the environment in which they act. Agents are concerned with a total, online, cost achievable in a limited time frame. In this context, agents must balance the coordinated *exploration* of the unknown environment and the *exploitation* of the known portion of the costs, in order to minimize the global utility. This model was originally called *Distributed Coordination of Exploration and Exploitation (DCEE)* (Taylor, Jain, Jin, Yokoo, & Tambe, 2010).

6.3.1 DEFINITION

The P-DCOP model for agents with partial knowledge is described by extending the P-DCOP model introduced in Section 6.1, as follows: $\langle \mathbf{A}, \mathbf{X}, \mathbf{D}, \mathbf{F}, \alpha, \mathcal{I}, \Omega, \mathbf{P}, \mathcal{E}, \mathcal{U}, T \rangle$, where $T > 0$ is a finite time horizon characterizing the time within which the agents can exploit the unknown cost functions and explore the search space. The goal in such a P-DCOP problem is to find a set of complete assignments $\vec{\sigma}^* = [\sigma_1^*, \ldots, \sigma_T^*]$ that minimizes the utility of the cumulative cost within the finite time horizon T:

$$\vec{\sigma}^* := \arg\min_{\sigma^1, \ldots, \sigma^T}/\max \ \mathcal{E}\left[\sum_{t=0}^{T} \sum_{f_i \in \mathbf{F}} \mathcal{U}\big(f_i(\sigma_{\mathbf{x}^i \setminus \mathcal{I}}^t)\big)\right] \tag{10}$$

where $\sigma^t \in \Sigma$ denotes a solution at time step t. In other words, agents have at most T time steps to modify the value of their decision variables, and solve T P-DCOP problems by acquiring more and more knowledge on the environment as the time unrolls.

6.3.2 ALGORITHMS

In a stochastic and unknown environment, the cost functions need to be learned online through interactions between the agents and their environment. Thus, the algorithms presented in this section are targeted to coordinate agents to solve a sequence of optimization problems in order to simultaneously reduce uncertainty about the local cost functions (exploration) and optimize the global objective (exploitation). In addition, the following algorithms are incomplete.

BE-Rebid (Taylor et al., 2010). The *Balanced Exploration Rebid (BE-Rebid)* is a synchronous, search-based algorithm that solves P-DCOPs with $\mathcal{I} = \emptyset$, and $\mathcal{U}$ and $\mathcal{E}$ are the identity functions. It extends MGM as it calculates and communicates its expected gain. The algorithm is introduced in the context of a *wireless sensor network* problem, where agents can perform small movements in order to enhance their communication capabilities, which are characterized by the distance between pairs of agents. Each agent can perform three actions: *stay* in the current position, *explore* another position, or *backtrack* to a previously explored position and halt movement. In each time step, BE-Rebid computes the expected cost of executing the *explore* or *backtrack* actions, assuming complete knowledge of the underlying distribution of the cost functions. Exploring is evaluated by using order statistics and is based on the cost of the best value found during exploration. Backtracking to

a known position results in a cost associated to the backtracked state for the remainder of the time steps (i.e., it *stays* in that state).

Following the region-optimal approaches presented in the context of classical DCOPs, Taylor et al. (2011) propose a version of the algorithm, called *BE-Rebid-2*, that allows pairs of agents to explore in a coordinated fashion. Interestingly, in such settings, the authors find that increasing coordination (measured by the number of agents that can execute a joint action) can decrease solution quality. This phenomenon is referred to as *team uncertainty penalty*. The worst case runtime and communication requirements of this algorithm are the same as those of MGM.

Heist (Stranders, Tran-Thanh, Fave, Rogers, & Jennings, 2012). *Heist* is a synchronous, inference-based algorithm that solves P-DCOPs with $\mathcal{I} = \emptyset$, $\mathcal{E}$ is the expectation function, and $\mathcal{U}$ is the identity function. It aims at minimizing the expected utility of the cumulative cost function $\mathbf{F}_g$, within the finite time horizon T. It does so by modifying a *Multi-Armed Bandit (MAB)* approach (Vermorel & Mohri, 2005) to a distributed scenario. A MAB is a slot machine with multiple arms, each of which yields a cost drawn from an unknown but fixed probability distribution. It trades exploration and exploitation by pulling the arms in order to minimize the cumulative cost over a finite horizon. To cope with the uncertain and stochastic nature of the cost functions, Heist models each cost function as a MAB, such that the joined assignment of the variables in the scope of the given cost function becomes an arm of that bandit. It seeks to minimize the expected cumulative optimization cost received over a finite time horizon by repeatedly pulling the MAB arms to select the joint action with the highest estimated *Upper Confidence Bound (UCB)* (Auer, Cesa-Bianchi, & Fischer, 2002) on the sum of the local gains received in a single time step. To do so, it employs a belief propagation algorithm, known as *Generalized Distributive Law (GDL)* (Aji & McEliece, 2000), in order to minimize the UCB in a decentralized fashion. Stranders et al. (2012) show that Heist enables agents to balance between exploration and exploitation, and derive optimal asymptotic bounds on the regret of the global cumulative cost attained.

The worst case runtime requirement of this algorithm is $O(\ell T d^l)$ as each agent computes the maximum marginal UCB for its variable assignment for each time step before the horizon. In terms of memory requirement, each Heist agent needs $O(d^l)$ space to store all value combinations of neighboring variables. In terms of communication requirement, the number of messages sent is $O(\ell T l)$, one to each neighbor and time step in each iteration, and the size of each message is $O(Td)$ as it contains the aggregated costs of all the agent's variable's values for each time step.

ICG-Max-Sum (Wu & Jennings, 2014). The *Iterative Constraint Generation Max-Sum (ICG-Max-Sum)* algorithm is a synchronous, inference-based algorithm that solves P-DCOPs with $\mathcal{I} \neq \emptyset$, $\mathcal{E}$ is the identity function, and $\mathcal{U}(f_i(\cdot))$ is the maximal regret function. The algorithm aims at minimizing the sum of maximal regrets for all the functions in $\mathbf{F}$. Furthermore, the horizon is $T = 1$. Thus, unlike the previous algorithms, IGC-Max-Sum does not attempt to learn the outcome of the cost functions. Its objective is to find *robust* solutions to the uncertain problem distributions; it does so by finding the solution that minimizes the maximum regret. The algorithm extends the *Iterative Constraint Generation (ICG)* method (Benders, 1962; Regan & Boutilier, 2010) to the decentralized case by decomposing the overall problem into a *master* problem and a *subproblem* that are iteratively solved until convergence. At each iteration, the resolutions of these problems are attained by using Max-Sum. The *master* problem solves a relaxation of the minimax regret goal, where only a subset of all possible joint beliefs is considered, attempting to minimize the loss for the worst case derived from the considered joint belief. Once it generates a solution, the *subproblem*

finds the maximally violated constraint associated to such a solution. This is referred to as the *witness* point, indicating that the current solution is not the best one in terms of the minimax regret. This point is added to the set of joint beliefs considered by the master problem, and the process is repeated until no new witness points can be found.

The worst case runtime requirement of this algorithm is $O(\ell|\mathcal{I}|d^l)$, which is dominated by the master problem, whose computation is exponential in the number of variables in the scope of the associated cost function for each belief and iteration of the algorithm. In terms of memory requirement, each ICG-Max-Sum agent needs $O(|\mathcal{I}|d^l)$ space (dominated by the master problem again) to store all value combinations of neighboring variables for each belief. In terms of communication requirement, the number of messages sent is the same as that of Max-Sum since two parallel iterations of Max-Sum is executed in each ICG-Max-Sum iteration. However, the size of each message is $O(|\mathcal{I}|d)$ as it contains the aggregated cost of all the agent's variable's value for each belief.

A variation of this algorithm that aims at minimizing the expected regret, rather than minimizing the maximum regret, was introduced by Le, Fioretto, Yeoh, Son, and Pontelli (2016).

7. Quantified DCOPs

The various extensions of the DCOP model discussed so far differ from each other in terms of agent behavior (deterministic vs. stochastic), agent knowledge (total vs. partial), environmental behavior (deterministic vs. stochastic), and environment evolution (static vs. dynamic). However, in terms of the agent teamwork, all of these models assume the agents are fully cooperative. Researchers have introduced the *Quantified DCOP (QDCOP)* model (Matsui, Matsuo, Silaghi, Hirayama, Yokoo, & Baba, 2010), which assumes a subset of agents to be adversarial, that is, the agents are *partially cooperative* or *competitive*.

7.1 Definition

The *Quantified DCOP (QDCOP)* model (Matsui et al., 2010) adapts the *Quantified Constraint Satisfaction Problem (QCSP)* (Benedetti, Lallouet, & Vautard, 2008) and *Quantified Distributed CSP (QDCSP)* (Baba, Iwasaki, Yokoo, Silaghi, Hirayama, & Matsui, 2010; Baba, Joe, Iwasaki, & Yokoo, 2011) models to DCOPs. In QCSPs and QDCSPs, all variables are associated to quantifiers and the constraints should be satisfied independently of the value taken by universally quantified variables. Analogously, in QDCOPs, existential ($\exists$) and universal ($\forall$) quantifiers are introduced to differentiate the cooperative agents from the adversarial ones.

A *QDCOP* has the form $Q(\mathbf{F}) \coloneqq q_0 x_0 \ldots q_n x_n$.[11] Q is a sequence of quantified variables, where each $q_i \in \{\exists, \forall\}$ quantifies the variable x_i. The goal of a QDCOP is to find a global optimal solution of the corresponding DCOP. However, a universally quantified variable is not coordinated nor assigned, as the result has to hold when it takes any value from its domain. In contrast, an existentially quantified variable takes exactly one value from its domain, as in (*cooperative*) DCOPs. Thus, the optimal solution of a QDCOP may be different from that of the corresponding DCOP. While a DCOP solution defines a single value, associated to its cost, a QDCOP defines *upper* and *lower bounds* to the optimal solution. In particular, the best choice in a QDCOP defines the smallest *lower bound*. In the worst case, the universally quantified variables can worsen the overall objective

---

11. In the original proposal, the set $\mathbf{F}$ is separated in a set of constraints C, representing relationships among variables, and a set of functions F, assigning values to each valid assignment.

as much as possible. Therefore, the worst case defines the smallest *upper bound*. While finding an optimal solution for a DCOP is NP-hard, solving a QDCOP is, in general, P-SPACE-hard (Benedetti et al., 2008; Lallouet, Lee, Mak, & Yip, 2015).

7.2 Algorithms

QDCOPs impose a rigid order on the variables, which reflects the correct order of evaluation of the quantifiers. Therefore, classical DCOP algorithms cannot be directly applied to solve QDCOPs. Matsui et al. (2010) proposed several variations of ADOPT to solve QDCOPs, which are all based on a DFS pseudo-tree ordering. To keep the ordering of the quantifiers unchanged, the pseudo-tree can be reshaped by applying extra null edges for each pair of nodes, if necessary.

All the algorithms presented here are complete, and are based on the intuition that universally quantified variables can be seen as adversarial virtual agents, whose goal is to minimize the overall objective. Following this intuition and the pseudo-tree modifications discussed above, pseudo-tree-based DCOP algorithms can be extended to solve QDCOP.

Min-max ADOPT (Matsui et al., 2010). *Min-max ADOPT* is an asynchronous, search-based algorithm that extends ADOPT to solve QDCOPs. It uses VALUE messages to communicate values of the variables, and COST messages to announce their costs, similar to ADOPT. Each agent, starting from the root of the pseudo-tree, assigns values to its variables and propagates them to its neighboring agents with lower priority. Upon receiving VALUE messages from all higher-priority neighbors, the agent updates its context and repeats the same process by choosing an assignment that minimizes its local cost. In Min-max ADOPT, the existentially quantified variables are used to compute the lower bound, while the universally quantified variables are used to compute the upper bound. This process is executed until the root agent detects that the upper bound is equal to the lower bound. This algorithm has a relatively simple structure and does not adopt any major pruning strategy.

Alpha-beta ADOPT (Matsui et al., 2010). *Alpha-beta ADOPT* is an asynchronous, search-based algorithm that extends Min-max ADOPT by adapting the *alpha-beta* search strategy, a common pruning strategy adopted in *game-tree search*. This strategy employs two boundary parameters, *alpha* and *beta*, representing the lower bound and the upper bound for each possible cost of an assignment, respectively. Alpha represents the lower bound, controlled by the universally quantified variables, while beta represents the upper bound, controlled by the existentially quantified variables. Lower bound and upper bound can be modified exclusively by universally quantified and existentially quantified variables, respectively. In Alpha-beta ADOPT, when an agent reports the cost value of the current partial assignment, its parent reduces the alpha/beta threshold accordingly. Thus, the new alpha/beta values are used to prune the search when an agent detects that the current assignment cannot be better than any other solution already evaluated. Alpha and beta values are obtained using a backtracking technique similar to how thresholds are obtained through backtracking in the original ADOPT.

Bi-threshold ADOPT (Matsui et al., 2010). *Bi-threshold ADOPT* extends ADOPT by employing two *backtracking thresholds* instead of one as in ADOPT. In ADOPT, each agent a_i maintains the threshold invariant $lb_i^* \leq t_i \leq ub_i^*$, where lb_i^* and ub_i^* are the smallest lower and upper bounds, respectively, of the agent over all of its values, and t_i is the threshold of the agent. In contrast, in Bi-threshold ADOPT, each agent maintains the threshold invariant $lb_i^* \leq t_i^\alpha \leq t_i^\beta \leq ub_i^*$, where t_i^α

Group	Problem	Model		
		AK	EB	EE
DISASTER MANAGEMENT & COORDINATION	*Disaster Evacuation*	T	D	S/D
	Coalition Formation	P	D	D
RADIO FREQUENCY ALLOCATION	*Cooperative Channel Assignment*	T	D	S/D
RECOMMENDATION SYSTEMS	*Group Recommendation*	T	D	S
SCHEDULING	*Distributed Meeting Scheduling*	T	D	S
	Water Allocation Scheduling	T	D	S
	Vessel Rotation Planning	T	D	S
	Patient Scheduling	T	D	S/D
SENSOR NETWORK	*Target Tracking*	T	D/S	D
	Robotic Network Optimization	P	D	D
	Mobile Sensor Team	P	S	D
	Sensor Sleep Scheduling	T	D	S
SERVICE-ORIENTED COMPUTING	*Application Component Placement*	T	D	S
	Server Allocation	T	D	D
SMART GRID AND SMART BUILDINGS	*Economic Dispatch*	T	D/S	S/D
	Power Supply Restoration	T	D	S
	Microgrid Islanding	T	D	S
	Prosumer Energy Trading	T	D/S	S
	Smart Building Devices Scheduling	T	D/S	S/D
SUPPLY CHAIN MANAGEMENT	*Supply Chain Formation*	T	D	S/D
TRAFFIC CONTROL	*Traffic Light Synchronization*	T/P	S	D

Table 5: DCOP Applications. AK = Agent Knowledge, it can be [T]otal or [P]artial; EB = Environment Behavior, it can be [D]eterministic or [S]tochastic; EE = Environment Evolution, it can be [S]tatic or [D]ynamic

is a lower bound on the threshold, similar to alpha in Alpha-beta ADOPT, and t_i^β is an upper bound on the threshold, similar to beta in Alpha-beta ADOPT.

8. DCOP Applications

DCOP models have been adopted to represent a wide range of MAS applications, thanks to their ability to capture essential and fundamental MAS aspects as well as the support for the development of general domain-independent algorithms. This section describes some of the most compelling applications as well as a general overview of their corresponding DCOP models. A comprehensive list of DCOP applications, categorized according to the DCOP classification of Table 2, is given in Table 5.

8.1 Disaster Management and Coordination Problems

Disaster management and coordination problems refer to how to efficiently and effectively respond to an emergency. In these scenarios, low-powered mobile devices that require limited bandwidth are often deployed and utilized to assist the disaster management process. Due to their decentralized nature, the DCOP approach fits naturally with this application. A description of several problems within this application domain is presented next.

8.1.1 DISASTER EVACUATION PROBLEMS

In a disaster scenario, moving evacuees to the closest refuge shelter can quickly overwhelm shelter capacities. A number of researchers have proposed a DCOP model for disaster evacuation, in which several groups of evacuees have to be led to available shelters (Carpenter, Dugan, Kopena, Lass, Naik, Nguyen, Sultanik, Modi, & Regli, 2007; Kopena, Sultanik, Lass, Nguyen, Dugan, Modi, & Regli, 2008; Lass, Kopena, Sultanik, Nguyen, Dugan, Modi, & Regli, 2008a; Lass, Regli, Kaplan, Mitkus, & Sim, 2008b; Kinoshita, Iizuka, & Iizuka, 2013). Group leaders can communicate via mobile devices to monitor and coordinate actions. Each group is represented by a DCOP agent managing variables that represent shelter allocations. Thus, the domain of each variable corresponds to the available shelters. Group sizes and shelter capacity, as well as additional group requirements (e.g., medical needs) and the distance of a group to shelters, are encoded as cost functions. Solving the DCOP ensures an assignment of all groups to shelters that minimizes overflow, such that groups receive the services they need and their travel distances are minimal.

8.1.2 COALITION FORMATION WITH SPATIAL AND TEMPORAL CONSTRAINTS PROBLEMS

In a *Coalition Formation with Spatial and Temporal Constraints (CFST)* problem (Ramchurn et al., 2010; Steinbauer & Kleiner, 2012; Pujol-Gonzalez, Cerquides, Farinelli, Meseguer, & Rodriguez-Aguilar, 2015), ambulance and fire brigade agents cooperate in order to react efficiently to an emergency scenario so as to rescue victims and extinguish fires located in different locations. Agents can travel from one location to another in a given time. Each task (i.e., rescuing a victim, extinguishing a fire) has a deadline, representing the time until which the victim will survive, and a workload, denoting the amount of time necessary to rescue the victim or put out the fire. The locations of the victims and the fires may be unknown to the agents, and need to be discovered at runtime, which requires agents to dynamically update the sequence of the tasks they will attempt, taking account of two main constraints: (1) Spatial constraints, which model where an agent can travel and at what time; and (2) Temporal constraints, which model task deadlines and completion times. Agents may also form coalitions to execute a given task faster or if the requirements of a given task cannot be met by a single agent. Hence, the agents' arrival times at each task need to be coordinated in order to form the desired coalition. The objective is to maximize the number of tasks to be completed.

A DCOP formalization for the CFST is described by Ramchurn et al. (2010), where ambulances and fire brigades are modeled as DCOP agents. Each agent controls a variable that encodes the current task that the agent will attempt, and whose domain represents task locations. Unary constraints restrict the set of reachable locations, according to distance from a destination and the victim's deadline. Agents' coalitions are defined as groups of agents traveling to the same location. Each task is associated to a utility, which encodes the success for a coalition to complete such task. The goal is to find an assignment of agents to tasks that maximizes the overall utility.

A variant of the CFST problem, called the *Law Enforcement Problem (LEP)*, was introduced by Amador, Okamoto, and Zivan (2014). Similar to the CFST problem, in a LEP, police officers need to execute a number of tasks, and may form coalitions to improve their response quality. Additionally, and differently from the CFST, new tasks can be revealed to the agents dynamically over time. Agents can choose to interrupt their current task to perform a new revealed task at the cost of a penalty that is proportional to the importance of the task being interrupted.

8.2 Radio Frequency Allocation Problems

The performance of a *wireless local area network (WLAN)* depends on the channel assignments among neighboring *access points (APs)*. Neighboring transmissions occurring in APs on the same channel or adjacent channels degrade network performance due to transmission interference. In dense urban areas, different APs may belong to different administrative domains, whose control is delegated to different entities. Thus, a distributed approach to the channel assignment is necessary.

8.2.1 Cooperative Channel Assignment Problems

In a cooperative channel assignment problem (Hollos, Karl, & Wolisz, 2004; Monteiro, Pellenz, Penna, Enembreck, Souza, & Pujolle, 2012a; Monteiro, Pujolle, Pellenz, Penna, Enembreck, & Demo Souza, 2012b), APs need to be configured in order to reduce the overall interference between simultaneous transmissions on neighboring channels. Monteiro et al. (2012a, 2012b) proposed a DCOP-based approach for cooperative channel assignment in WLANs where APs may belong to different administration entities. In the proposed model, each AP is represented by a DCOP agent, which controls a decision variable modeling a choice for the AP's channels. The signal-to-interference-and-noise ratio perceived by an AP or a client is modeled as a cost function, as the overall concurrent transmissions occurring in the same channel and in partially overlapped adjacent channels. The goal is to find an assignment of channels to APs that minimizes the total interference experienced in the WLAN.

Xie, Howitt, and Raja (2007) and Mir, Merghem-Boulahia, and Gaïti (2010) study dynamic solutions to the problem of allocating and utilizing the wireless network's available spectrum. In such a problem, the agents operate in a dynamic radio frequency environment that is composed of time-varying interference sources, which are periodically sampled and measured.

8.3 Recommendation Systems

Recommendation systems are tools that provide user-tailored information about items to users. These systems provide information that is tailored to the characteristics and preferences of the users.

8.3.1 Group Recommendation Problems

Like individual recommendations, group recommendations need to take into account the preferences of all group members and formulate a recommendation that suits the whole group. Lorenzi, dos Santos, Ferreira Jr, and Bazzan (2008) propose a DCOP-based travel package recommendation system for groups of users. The users in the group share a common goal (the travel package recommendation), and have individual preferences for each travel service (hotel, flight companies, tour operators, etc). In such a problem, the objective is to find a group recommendation that optimizes the users' preferences. The proposed DCOP solution is composed of two types of agents: user agents

and recommender agents. Each user agent controls a decision variable that models that user's travel choices, while each recommender agent controls a decision variable that models a travel service supplier's recommendations. User travel preferences are modeled via unary constraints on the user agents' decision variables. Binary constraints between user agents in a group and their associated recommender agent ensure that each user's choice in the group is compatible to the recommendation of the recommender agent. The goal is to find the best recommendation for the entire group.

8.4 Scheduling Problems

Scheduling problems are an important class of problems that have been long studied in the area of constraint programming and operations research (Giffler & Thompson, 1960; Solomon, 1987; Minton, Johnston, Philips, & Laird, 1992; Van Hentenryck & Michel, 2009). In such problems, time schedules are to be associated to resource usage. The problem is made particularly difficult when the scheduling process needs to be coordinated in a distributed manner across several entities. In such a context, many scheduling problems can be naturally mapped to DCOPs.

8.4.1 DISTRIBUTED MEETING SCHEDULING PROBLEMS

The distributed meeting scheduling problem captures generic scheduling problems where one wishes to schedule a set of events within a time range (Jennings & Jackson, 1995; Garrido & Sycara, 1996). Each event is defined by: *(i)* the resources required for it to be completed, *(ii)* the time required for it to be completed, within which it holds the required resources, and *(iii)* the cost of using such resources at a given time. A scheduling conflict occurs if two events with at least one common resource are scheduled in overlapping time slots. The goal is to maximize the utilities over all the resources, defined as the net gain between the opportunity benefit and opportunity cost of scheduling various events.

Maheswaran, Tambe, Bowring, Pearce, and Varakantham (2004b) discuss three possible DCOP formulations for this problem: *Time slots as variables (TSAV)*, *events as variables (EAV)*, and *private events as variables (PEAV)*. We describe the EAV formulation and refer the reader to the original article for the other two formulations and additional details. In the EAV formulation, events are considered as decision variables. Each variable can take on a value from the time slot range that is sufficiently early to schedule the required resources for the required amount of time, or zero to denote that an event is not scheduled. If a variable takes on a non-zero value, then all its required resources cannot be assigned to any other overlapping event.

8.4.2 WATER ALLOCATION SCHEDULING PROBLEMS

The management of water resources in large-scale systems is often associated with multiple institutionally-independent decision makers, which may represent different and conflicting interests, such as flood prevention, hydropower production, and water supply (Giuliani, Castelletti, Amigoni, & Cai, 2014). The aim of such problems is to find an efficient use of water allocation and distribution according to the different users' interests.

Giuliani et al. (2014) formalize a regulatory mechanism in water management as a DCOP. The model involves several *active* human agents and *passive* ecological agents. Each agent is associated with an objective function that it seeks to maximize. Active agents make decisions about the amount of water to divert from the river or to be released from a dam in order to maximize their corresponding objective functions. Passive agents, on the other hand, represent ecological interests

through their associated objective functions and do not make decisions. The agents model different water supplies for cities and agricultural districts, hydropower productions, and ecological preservation. The goal is to optimize the agents' objective functions, satisfying hard (physical) constraints and maximizing the soft (normative) constraints, which aim at protecting the interests of the passive agents. A solution to such problem, which makes use of a multi-objective DCOP formalization, is presented by Amigoni, Castelletti, and Giuliani (2015).

8.4.3 Vessel Rotation Planning Problem

In a large port, vessels are scheduled to visit different terminals for loading and unloading operations. The sequence of terminals visited by a vessel is called the *vessel rotation*. Due to the different nature of the terminal visits (e.g., loading, unloading containers, different inland shipping activities) there are often dependencies between activities performed at the terminals. Additionally, vessel operators have their own preferences on when to visit a particular terminal. The *vessel rotation planning problem (VRPP)* (Li, Negenborn, & Lodewijks, 2016) describes the problem of assigning rotations to vessels, consisting of sequences of visits to terminals and arrival and departure times of vessels to terminals in a port area, while satisfying the terminals' visits dependencies and taking into account operators' preferences. Due to the geographically distributed nature of the resources (vessels, terminals, loading facilities) and the distributed coordination process undertaken by vessels operators, this problem fits naturally in the DCOP framework.

Li et al. (2016) proposed a DCOP model for the VRPP, where vessels and terminals are modeled as DCOP agents. The problem's variables are the time slot for which a vessel i is at terminal j (x_{ij}), the arrival and departure times of vessel i at terminal j (a_{ij} and d_{ij}, respectively), and the number of time steps a vessel i waits at terminal j (w_{ij}). A vessel agent i governs variables x_{ij}, a_{ij}, d_{ij}, and w_{ij}, for each terminal j that it can visit, while a terminal agent controls auxiliary variables y_{ij} for each vessel i that can visit j, which has the same value of x_{ij} and is used by terminal agents to represent the terminal capacities. The problem constraints represent the preferences of vessels of being at a given terminal in a given time, the penalty occurring for a waiting vessel, and the time windows during which a vessel can visit a terminal.

8.4.4 Patient Scheduling Problems

Medical appointments scheduling problems are related to meeting scheduling problems, as they need to associate patients to resources (e.g., doctors, medical machinery) and times, but they require different types of constraints. Patients may require several services from different departments within the same hospital or in multiple hospitals. In general, the objective is to minimize the patient treatment waiting time under limited resource conditions, as well as to ensure efficient resource usage, taking into account patient preferences.

Hannebauer and Müller (2001) formulate the problem of scheduling patients to diagnostic units in an hospital as a DCOP, where appointments are modeled as variables, whose domains describe times, durations, and locations. The constraints of the problem model the schedule feasibility, the patient preferences over hospitalization times, the workplace constraints, which restricts the types of appointment for a given workplace, and diagnostic unit constraints, which model resource usage.

Billiau, Chang, Ghose, and Miller (2012b, 2012c) propose a Dynamic DCOP model for a radiotherapy patients scheduling problem. In this problem, each agent represents a patient, and it controls variables that represent private information (e.g., type of tumor, number of radiation doses

per day, the use of chemotherapy) and public information (e.g., current schedule of the radiotherapy machine). The constraints of the problem model the duration of each daily treatment, as well as tumor-specific treatment restrictions. The problem objective considers patient waiting times to receive their treatment, patient priorities (based on tumor aggressiveness), and patient preferences.

8.5 Sensor Network Problems

Sensor networks typically consist of a large number of inexpensive and autonomous sensor nodes, constrained by a limited communication range and battery life. These networks have been deployed for environmental sensing (temperature, humidity, etc.), military applications (e.g., battlefield surveillance), and target tracking (Akyildiz, Su, Sankarasubramaniam, & Cayirci, 2002). When deploying sensor networks, it may not be possible to pre-determine the position of each sensor node. The distributed nature of the problem and the presence of several communication and sensing constraints create a natural fit for DCOPs to solve a wide range of related applications.

8.5.1 TARGET TRACKING PROBLEMS

In a target tracking application (Zhang et al., 2005; Matsui & Matsuo, 2008; Jain, Taylor, Tambe, & Yokoo, 2009; Ota, Matsui, & Matsuo, 2009; Stranders, Farinelli, Rogers, & Jennings, 2009; Semnani & Basir, 2013), a collection of small Doppler sensors are scattered in an area to detect possible moving targets in the region. Each sensor is battery-powered, can communicate with one another through radio communication, and can scan and detect an object within a fixed range. Communication incurs an energy cost. Thus, to save energy, a sensor may turn itself off. Multiple sensors may be necessary to detect a single target with high accuracy. The overall objective is to maximize the number of targets detected, as quickly as possible and, at the same time, preserve energy so as to prolong the system's lifetime

Zhang et al. (2005) model a simplified version of the above problem as a weighted graph coloring problem, where the total weight of violated constraints needs to be minimized. A node corresponds to a sensor, an edge between two nodes represents the constraint of a shared region between agents, and the weight captures the importance of the common region. The size of the common region reflects the amount of energy loss when two sensors scan the shared region at the same time. Each color corresponds to a time slot in which a sector is scanned. A node must have at least one color so that the corresponding sector is scanned at least once. This graph coloring problem is mapped to a DCOP, where agents represent nodes, agent's variables represent the agents decision on their color, and cost functions represent the graph edges.

Semnani and Basir (2013) use a hierarchical DCOP approach to scale to larger problems. The authors partition the original problem into n local regions, and use n DCOPs to solve the smaller subproblems. Their solutions are then combined in a hierarchical approach, solved by a DCOP that encompasses variables and constraints shared among the connected regions of the lower hierarchy DCOPs.

8.5.2 ROBOTIC NETWORK OPTIMIZATION PROBLEMS

The robotic network optimization problem describes a sensor network problem where sensors are placed on top of robots that have limited movement capability. In such a problem, robots can make small movements to optimize the wireless connectivity with their neighbors, without affecting the network topology (Choxi & Modi, 2007; Jain et al., 2009).

Jain et al. (2009) proposed a DCOP formulation where each robot is represented by an agent. Each agent controls one variable describing the decision on the robots' possible movements. Thus, the variables' domains consist of the valid positions the agent can move to. The cost functions of the problem model the power loss (or gain, depending on the optimization criteria) of the wireless link from a transmitter and a receiver robot, and depend on their positions. Radio communication in wireless sensor networks have a predictable signal strength loss that is roughly inversely proportional to the square of the distance between transmitter and receiver. However, radio wave interference is very difficult to predict (Molisch, 2012). Thus, Jain et al. (2009) use a P-DCOP-based approach with partial agent knowledge to capture the robot's partial knowledge on its cost functions, and to balance *exploration* of the unknown costs and *exploitation* of the known portion of the costs.

8.5.3 Mobile Sensor Team Problem

The *Mobile Sensor Team (MST)* problem is similar to the target tracking problem with the difference that agents are capable of moving autonomously within the environment and that time is modeled explicitly as a discrete sequence of time steps. In an MST, agents are placed on a grid. For an agent a_i, cur_pos_i denotes the agent's current position; SR_i denotes the agent's perception *sensing range*, which determines the coverage range within which an agent can detect targets; MR_i denotes the agent's mobility range, which defines the maximum distance that the agent can move within a single time step; and $cred_i$ denotes the agent's credibility, which reflects the likelihood of the correctness of the detected targets. The targets are defined implicitly through an *environmental requirement (ER)* function, which defines, for each point in the space, the minimum *joint credibility* value (the sum of the credibility variables) required for that point to be sensed. In such a representation, targets are points p with $ER(p) > 0$. Given a set of agents SR_p whose sensing range covers a target p, the *remaining coverage requirement* of p is the environmental requirement diminished by the joint credibility of the agents currently covering p: $Cur\_REQ(p) = \max\{0, ER(p) \ominus \sum_{a_i \in SR_p} cred_i\}$, where $\ominus : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ is an operator that defines how the environmental requirement decreases by the joint credibility. The goal of the agents is to find positions that minimize the values of Cur_REQ for all targets.

MST problems are modeled through a subclass of Dynamic DCOPs, named *DCOP_MST* (Zivan, Glinton, & Sycara, 2009; Yedidsion, Zivan, & Farinelli, 2014; Yedidsion & Zivan, 2014). Each agent a_i controls one variable x_i representing its position, and whose domain contains all locations within MR_i of cur_pos_i. Thus, the domains are updated each time the agent moves.[12] The constraint C_p of a target p involves exclusively those agents a_i whose variable's domain includes a location within SR_i of p. Thus, at each time step, both domains and constraints may change. As a consequence, the constraint graph changes as well – the neighbors of each agent have to be updated at each time step. Finally, in a DCOP_MST two agents are neighbors if their sensing areas overlap.

8.5.4 Sensor Sleep Scheduling Problem

Wireless sensor nodes are equipped with a radio, which can be used to communicate with neighboring nodes, and a limited power source. These sensor nodes are often deployed in inaccessible terrains, thus, replacing their power sources may not be possible. The wireless sensor sleeping

---

12. An alternative representation is that of modeling the domain of each variable as the entire grid, and to constrain the agent variable, at each agent move, in order to hide those points lying outside MR_i.

scheduling problem aims at switching on/off a particular sensor node component (such as the sensor or the radio) for a certain period of time, so to ensure power conservation, maximizing the lifetime of the sensor network.

Chachra and Marefat (2006) proposed a DCOP model for this problem, where each sensor is an agent whose variables denote its status (on or off) for each time step. Hard constraints are employed to enforced that if a sensor is on, then all its neighbors should be off, and that sensors cannot stay on for two consecutive time steps. The overall objective is to minimize the delay induced in the network.

A similar problem is solved by Stranders et al. (2009), where sensors are also able to harvest energy from the environment (e.g. using a photo-voltaic cell or vibration-harvesting microgenerators). In such a context, the goal is to find a schedule that maximizes the probability of detecting events while maintaining energy-neutral operations (that is, exhibit an indefinite lifetime for each of the agents).

8.6 Service-Oriented Computing Problems

The service-oriented computing paradigm is one that relies on sharing resources over a network, focusing on maximizing the effectiveness of the shared resources, which are used by multiple applications. Efficient solutions with optimal use of resources are crucial in this paradigm and have a wide industrial impact (Moreno-Vozmediano, Montero, & Llorente, 2013). The distributed nature of the resources and the privacy concerns arising when different clouds are involved in the deployment, makes DCOP appealing to solve a range of problems in this paradigm (Mejias & Roy, 2010).

8.6.1 APPLICATION COMPONENT PLACEMENT PROBLEMS

An *Application Component Placement (ACP)* problem is defined over a network of servers offering storage devices with various capabilities, and component-based application with requirements for processing, communication, and storage (Jin, Cao, & Li, 2011; Li, Wang, Ding, & Li, 2014). The ACP problem is a problem of deciding which server to assign to each application component. The component-based application is described by a set of characteristics that establish their requirements in terms of hardware (e.g., CPU speed, storage capacity) as well as constraints between components of the same application (e.g., minimum bandwidth, secure communication channel requirement). When the APC involves deployment on multiple clouds data privacy must be preserved. Additionally, in cloud environments computing resources are shared by many applications and the infrastructure is dynamically changing, making centralized solutions unfeasible.

Jin et al. (2011) proposed a DCOP model for the ACP problem where servers bid for a component to host, with an emphasis that is proportional to the affinity of the server characteristics and the component hardware and software requirements. Each server is modeled by an agent, which controls a decision variable representing the server bids. Thus, the domain of each variable is the set of possible components that may be deployed on the server. Unary functions express utility for each component. Hard constraints are employed to ensure that each component is deployed exactly on a single server, and that two components are placed between servers satisfying the required communication bandwidth. The objective is to find a feasible assignment of component to servers that maximizes the utilities.

8.6.2 Server Allocation Problems

Services-oriented middleware networks are composed of entities connected within a physical network. These entities can both provide and require multiple services. In turn, each service can be provided by multiple servers and can serve multiple clients. A service request from a given client, takes into account various *Quality of Service (QoS)* parameters (e.g., service response time, service completion time). When a client generates a service request, it can be satisfied by any of the servers offering such requests. The server allocation problem is a problem of selecting servers to allocate services, ensuring maximum social welfare, while meeting the QoS requirements of all clients.

Choudhury, Dey, Dutta, and Choudhury (2014) presented a DCOP model for this problem where agents correspond to network entities, variables correspond to services, and their values are either 1, if the associated agent is willing to provide/forward the service, or 0, otherwise. Clients' service requests are mapped to servers' service offers, accounting for the delays that occur when traversing between intermediate nodes using a routing multicast protocol (Yan, Lee, Shen, & Qiao, 2013). Moreover, in order to provide a service, all requested QoS requirements need to be satisfied. The utility associated to each variable is the combination of the utility for such a node when acting as a service consumer, a service provider, and a service forwarder, and depends on several parameters, such as available GPU cycles, battery power, memory, and bandwidth. The problem may change dynamically when a new request is made, or when a new service is offered or released, and as such can be modeled as a Dynamic DCOP.

8.7 Smart Grid and Smart Homes Problems

The smart grid is a vision of the future electricity grid (also called power network) that uses data analytics and decision making to improve the efficiency and reliability of energy production and distribution. The development of smart grids poses several challenges: *(i)* How to deal with the increasing power network utilization due to growth of loads, such as *electric vehicles (EVs)* and heat pumps; *(ii)* How to efficiently integrate a diverse range of energy sources, including renewable generators, into the power network; and *(iii)* How to deal with the uncertainty in the equipment as well as in the participation of consumers through demand-side technologies. Due to the distributed and dynamic nature of loads and generators participating in the power network, agent-based decentralized autonomous control of smaller distributed microgrids is a very compelling solution (Davidson, Dolan, McArthur, & Ault, 2009; Ramchurn, Vytelingum, Rogers, & Jennings, 2012). In particular, several agent-based decentralized optimization solutions have been explored to deliver this vision (Kumar, Faltings, & Petcu, 2009; Matsui & Matsuo, 2011; Miller, Ramchurn, & Rogers, 2012; Jain, Ranade, Gupta, & Pontelli, 2012). The following is a list of the most prominent DCOP approaches for smart grid and smart homes applications.

8.7.1 Economic Dispatch Problems

Economic Dispatch (ED) is the problem of coordinating the various settings of the power generators in order to meet the power loads with the lowest cost possible, while satisfying the physical power network constraints (Wood & Wollenberg, 2012). Researchers have cast this problem as a DCOP (Miller et al., 2012; Jain et al., 2012; Gupta, Jain, Yeoh, Ranade, & Pontelli, 2013a; Athanasiadis, Kockar, & McArthur, 2013), considering a network of nodes (agents), each of which relays power to other nodes, but can also contain a combination of generators and loads. Generators are distributed across nodes, and are represented through variables whose domain describe a certain set of discrete

power outputs. The power lines connecting nodes of the networks are also associated to DCOP variables, each of which has a thermal capacity describing the maximum power that it can safely carry. The DCOP objective function describes a particular optimization criteria, such as minimizing the carbon emissions of generators within the network, as well as the imposed load and network constraints. In particular, the constraints ensure that the overall demand and supply are in balance and that the thermal capacity constraints of the power lines are satisfied.

A version of the ED problem with a planning horizon has been proposed by Fioretto, Yeoh, Pontelli, Ma, and Ranade (2017b). This formalization extends the ED model described above by capturing the physical restrictions of transmission lines, power loads, and power generators arising when deploying a sequence of solutions (set-points for generators and loads) over time. In particular, it models the maximum incremental power that can be supplied or reduced in one time step to each power generator, which depends on the mechanical characteristics of the generators. This problem has been cast as a Dynamic DCOP, where, in addition to the variables and constraints of the classic ED problem, a set of constraints is introduced between each two consecutive time steps to limit the maximum variation of the generators' output.

8.7.2 POWER SUPPLY RESTORATION PROBLEMS

After (multiple) line failures, a power network must be reconfigured to ensure restoration of power supply. A power network distribution is a network of power lines connected by *switching devices (SDs)* and fed by *circuit breakers (CBs)*. SDs are analogous to sinks (transformer stations), while CBs are analogous to power sources. Both of these devices can operate in two states: open or closed. Closed SDs consume some power and forward the rest of it on other lines. Open SDs stop power flow. CBs feed the network when they are closed. The configuration of the devices' state is such that energy flow traversing CBs takes the form of a (*feeder*) *tree*, and that no SD is powered by more than a single power line. Flow conservation and transmission line capacity constraints must be enforced. The power supply restoration problem is the problem of finding a configuration that ensures power restoration for the maximum number of sinks affected by the line failures.

Researchers have proposed a DCOP formulation for this problem (Kumar et al., 2009; Agrawal, Kumar, & Varakantham, 2015). In such a framework, each node of the distribution network is controlled by an agent that owns all variables and constraints corresponding to that node. Two DCOP variables are associated to each network node: A load variable and a direction variable. Load variables model the amount of incoming flow for sink nodes, and the number of sinks fed for power source nodes. Direction variables model all the possibilities of feeding a node, as the set of possible configurations in which its neighboring nodes can forward power to it. The acyclicity of the power flow and the flow conservation are modeled as constraints. The former restricts the power path to be a tree as well as defines the optimization criterion. The latter enforces Kirchhoff's law, that the amount of incoming power flow to the node i must equal the sum of power consumed at i and the amount of power forwarded to other nodes.

8.7.3 MICROGRID ISLANDING PROBLEMS

A microgrid islanding problem is the problem of creating islands (i.e., clusters of generator units and loads able to operate without external energy supply) in response to major power outages and blackouts. Gupta, Yeoh, Pontelli, Jain, and Ranade (2013b) formalized this problem as a DCOP where agents represent nodes in the network and each agent has its own power generation and

power consumption capabilities. The variables of the DCOP represent the amount of power that an agent generates and consumes, as well as transmission line flows and the switch status between network nodes. The flow variables are constrained by their maximum transmission line capacities, while switches are modeled as binary variables that can be turned on or off. Flow conservation are modeled as constraints to enforce Kirchhoff's law. The goal is to find a switching configuration that minimizes the unserved load of the system.

8.7.4 PROSUMER ENERGY TRADING PROBLEMS

In its more general form, a smart grid is populated by prosumers capable of both generating and consuming resources. The prosumer energy trading problem aims at setting market-based prices for prosumers to directly trade energy over the smart grid, while taking account of the power network constraints. This problem has been cast as an optimization problem, called the energy allocation problem, where, given a graph with nodes representing prosumers and edges describing transmission lines connecting adjacent prosumers, the goal is to find an allocation that maximizes the benefits of all the prosumers while satisfying the capacity constraints of the power network.

Cerquides, Picard, and Rodríguez-Aguilar (2015) proposed a DCOP formalization for this problem, where each prosumer is modeled as an agent. Variables are associated with edges of the energy trading network (i,j) and describe the number of units of energy that prosumer i sells to/buys from prosumer j. Thus, two variables (i,j) and (j,i) are associated to each edge of the network. For each prosumer, an energy balance constraint models the utility of a given instantiation of its offers as the sum of the offers associated to the energy traded with each of its neighbor. Line capacity and flow conservation constraints ensure that the energy traded along the transmission lines is within their maximum capacity and is consistent with Kirchhoff's law.

8.7.5 SMART BUILDING DEVICES SCHEDULING PROBLEMS

A *smart building* is a residential or commercial building that is partially automated through the introduction of smart devices (e.g., smart thermostats, circulator heating, washing machines). Additionally, a range of smart plugs allow users to control remotely the activity of the devices connected to them. Therefore, smart device scheduling can be executed by householders without the control of a centralized authority. The distributed nature of smart devices within a smart building, and of smart buildings within a neighborhood, as well as data privacy concerns, make this domain suitable to DCOP solutions.

Within a smart building, the *Smart Environment Configuration Problem (SECP)* proposed by Rust, Picard, and Ramparany (2016) is the problem of coordinating several smart devices (light bulbs, roller shutters, luminosity sensors, etc.) whose actions affect the building environment, with the goal of reaching a desired goal state (e.g., a given luminosity level for a room). This resource allocation problem is cast to a DCOP in which devices are described via DCOP agents, each of which controls a single variable describing the devices' action. The domain of the variables model the possible device's actions. Finally, the impact of the device's action onto the building environment is captured by a set of soft constraints. The goal is that of finding an assignment of actions for the building's devices that satisfies the given goals while minimizing some cost function (e.g., the energy consumption of each device).

Within a smart city, the *Smart Home Device Scheduling Problem (SHDS)* proposed by Fioretto, Yeoh, and Pontelli (2017a) is the problem of coordinating the schedule of several smart homes' de-

vices so to minimize the aggregated energy peaks as well as to minimize the users' energy bill cost, when adopting a real-time energy price schema. In the SHDS problem, each smart home controls a set of smart sensors and smart devices. Within each smart home, users establish scheduling rules defining goal states upon certain home's properties (e.g., reaching a certain room temperature by a given time of the day, or charging the battery of an electric vehicle of at least some amount during the night). Each device's action has an associated energy consumption and a cost (which depends on the time of the day in which the action is executed and on its required power). Each user in a smart home attempts to find a schedule for its smart devices that satisfies its scheduling rules while minimizing the energy cost. Additionally, multiple smart homes coordinate their devices' schedules so to minimize the daily energy peak consumption. Fioretto et. al map the SHDS problem as a DCOP where each agent models a smart home. Agents control multiple variables, each representing a smart device (a sensor or an actuator). Similar to the SECP problem, variables' domains model the possible actions of a device, and constraints capture the actions' costs, energy consumption, and user preferences. Additionally, a set of hard constraints is used to model the user's temporal goals. A dataset for the SHDS problem has been recently released (Kluegel, Iqbal, Fioretto, Yeoh, & Pontelli, 2017). Finally, Tabakhi, Le, Fioretto, and Yeoh (2017) investigate how to elicit preferences associated to the users' temporal goals in an extension of the SHDS problem.

8.8 Supply Chain Management Problems

The management of large businesses involves the management of the flow of goods from suppliers to customers. This flow of goods is called a supply chain. Supply chains have to be carefully managed to ensure that a sufficient quantity of raw material is available at factories for production and a sufficient quantity of processed goods is available at stores for consumers to buy. Additionally, since goods can be purchased from different producers and sold to different consumers, there is also the need to consider how much to buy/sell the goods and who to buy/sell the goods to. In such an environment, information, decision making and control are inherently decentralized.

8.8.1 SUPPLY CHAIN FORMATION PROBLEMS

A supply chain formation problem is the process of determining the participants in a supply chain, who will exchange what, with whom, and the terms of the exchange. Several DCOP-based approaches for this problem have been proposed in the literature (Gaudreault, Frayret, & Pesant, 2009; Penya-Alba, Cerquides, Rodriguez-Aguilar, & Vinyals, 2012; Penya-Alba, Vinyals, Cerquides, & Rodriguez-Aguilar, 2012; Winsper & Chli, 2013; Penya-Alba, Vinyals, Cerquides, & Rodriguez-Aguilar, 2014). They rely on the notion of a *Task Dependency Network (TDN)*, a graph-based representation to capture dependencies among production processes, introduced by Walsh and Wellman (2000). A TDN is a bipartite directed acyclic graph representing exchanges, where nodes in the graph correspond to producers and consumers and edges in the graph correspond to feasible exchanges of goods between the producers and consumers. A path from potential producers and consumers defines a feasible supply chain configuration and the goal is to find the feasible configuration that optimizes a particular cost function. A DCOP encoding for this problem models producers and consumers as agents, each of which controls a variable describing the agent's decision regarding whom to buy from/sell to as well as its associated quantities and costs. Cost functions are associated with each edge of the TDN encoding the willingness of the producer and consumer to trade with each other.

A dynamic version of the above formalization has been investigated by Chli and Winsper (2015). This model allows for the entry and departure of producers and consumers as well as changes in properties of the problem (e.g., prices of goods, production capacity of producers, and consumption requirements of consumers).

8.9 Traffic Flow Control Problems

A challenge for the increase of transportation demand is to enforce traffic flow control using the existing infrastructure, such as traffic lights, loop detectors, and cameras. Coordinating the actions of these individual devices aims at smoothing the traffic flow at the network level. Such coordinated actions often generates coherent traffic control plans faster and more accurately compared to those of a human traffic operator (Van Katwijk, De Schutter, & Hellendoorn, 2009). Due to the distributed nature of such devices, multi-agent solutions are particularly suitable for this class of problems.

8.9.1 TRAFFIC LIGHT SYNCHRONIZATION PROBLEM

This problem is the problem of finding a synchronization schema for the traffic lights in adjacent intersections that creates green waves, which are waves of vehicles that are traveling at a given speed and are able to cross multiple intersections without stopping at red lights.

De Oliveira, Bazzan, and Lesser (2005) and Junges and Bazzan (2008) model this problem as a DCOP, where agents represent traffic lights, each controlling one variable that models the coordination direction for the associated traffic signal. Thus, the domain of the variables is given by two possible directions of coordination (north-south/south-north, east-west/west-east). Conflicts that arise when two neighboring traffic signals choose different directions are modeled as constraints. Cost functions are defined to model the number of incoming vehicles in a junction, and the costs are influenced by whether two adjacent agents agree on a direction of coordination or not. The goal is to minimize the global cost. Due to the dynamic nature of the problem, the proposed algorithms fall under the umbrella of Dynamic DCOP algorithms.

Pham, Tawfik, and Taylor (2013) and Brys, Pham, and Taylor (2014) proposed a Probabilistic DCOP-based approach with partial agent knowledge to solve the traffic light synchronization problem in the context where agents have partial information about their cost functions. In particular, the authors argue that traffic patterns may vary during time and, thus, the agents should learn them to update their cost functions. In this context, agents are given a limited amount of time to explore different signal duration intervals, whose associated costs are learned by evaluating the average travel time of the first 100 cars traveling across the agent's traffic light.

9. Analyses and Perspectives on DCOPs

DCOPs have emerged as a popular formalism for distributed reasoning and coordination in multi-agent systems. It provides an elegant modeling framework, which offers flexibility in both the agents' reasoning and coordination strategies. Its ability to support the notions of *preferences* and *constraints* makes it suitable to model a variety of multi-agent optimization problems. Preferences are a central concept of decision making and arise in a multitude of contexts in multi-agent systems. They are fundamental for the analysis of human choice behavior, and allow agents to express their inclinations through specific actions and behaviors. Constraints have been long studied in centralized systems (Rossi et al., 2006) and have been proved especially practical and efficient for

modeling and solving resource allocation and scheduling problems. They are naturally handled within DCOPs and offer a flexible and effective mean to model a variety of complex problems. In addition, DCOPs support several aspects that are crucial in multi-agent systems, such as agent privacy, autonomy in reasoning, and cooperation.

The classical DCOP notion is unable to capture important aspects of the problem related with the environment characteristics, such as partial observability, environment evolution, and uncertainty. Therefore, several DCOP model extensions have been recently presented. Each DCOP model imposes distinctive algorithmic requirements, which concern both the agent's reasoning and cooperation aspects. These requirements, in turn, are strictly related to the characteristics of the problem domain. Due to the performance variability imposed by such requirements, an appropriate selection of the DCOP model and algorithm is essential to obtain desirable performances in realistic application domains.

9.1 Comparative Overview of DCOP Models

Classical DCOPs can be used to represent a wide range of MAS applications where agents in a team need to work cooperatively to achieve a single goal in a static, deterministic, and fully observable environment. Exploring the domain structural properties, as well as understanding the requirements of the problem designer, is crucial to design and apply effective DCOP algorithms. A discussion on how to choose an appropriate algorithm based on the characteristic of the application at hand is provided in Section 4.4.

Asymmetric DCOPs are suitable when constrained agents incur different costs for a joint action, which arise especially in scenarios where privacy is a particular concern, where agents cannot reveal to the other agents the costs associated to their putative actions. Examples of problems that can be suitably modeled as Asymmetric DCOPs are resource allocation problems where agents incur in distinct costs from using the same resource, and where their preferences and constraints regarding usage time slots and durations are expected to be different. Asymmetric DCOPs are particularly attractive to model those domains that can be represented as *graphical games* (Kearns, Littman, & Singh, 2001), and where constraint reasoning could be actively used to exploit the problem structure. In a graphical game, the costs of each agent are affected exclusively by its neighboring agents. It is important to note that, even though the Asymmetric DCOP model bears similarities with many game-theoretic approaches, these two models are fundamentally different. While game-theoretic agents are self-interested, and their non-cooperative actions lead to a desirable global target, Asymmetric DCOP agents are cooperative and seek to minimize the global cost even at the expense of larger local costs.

Multi-Objective DCOPs are tailored to represent those classes of distributed problems that cannot be modeled with a single optimization function. The observations above on classical DCOPs also hold for Multi-Objective DCOPs and, in addition, agents cooperate to optimize multiple objectives. Numerous MAS applications fall in the category of multi-objective optimization. One example is that of *disaster management and coordination problems*, where agents coordinate to effectively respond to an emergency scenario (see Section 8.1). Due to memory requirements, which are proportional to the number of objectives and the size of the Pareto set, incomplete strategies seem particularly promising for this research area.

Dynamic DCOPs capture the dynamic behavior of the evolving environment in which agents act. Dynamic environments play a fundamental role in real-wold MAS applications. Virtually all

complex MAS applications involve dynamic situations, which may restructure the network topology due to agent movements, or bring additional information to the problem being solved. For example, in a search and rescue operation during *disaster management*, as the environment evolves over time, new information becomes available about civilians to be rescued and new agencies may arrive at any time to help conduct rescue operations (see Section 8.1). In a *smart grid domain*, real-time pricing is commonly enforced. Thus, agent preferences need to be adapted over time, while energy costs are updated (see Section 8.7). Dynamic DCOPs are therefore a modern area that presents an exciting field for groundbreaking research.

Probabilistic DCOPs extend the classical DCOP model to include the capability of handling uncertain events, allowing DCOP agents to handle a wider range of applications. In particular, Probabilistic DCOPs are suitable to capture those applications characterized by a static environment evolution with exogenous uncertain events (e.g., when the actions of agents on the environment can have different outcomes, based on external, uncertain factors) and, yet, agents have total knowledge of their own actions and of the observable environment and act in a fully cooperative context. The domain of *multi-agent task planning and scheduling* encompasses diverse problems that require complex models and robust solutions under uncertainty (see Section 8.4). The Probabilistic DCOP model for agents with partial knowledge is suitable to model those applications where agents have no prior knowledge of how the environment reacts to some of the actions. In such a model, the agents are aware of their own actions, which are performed deterministically. However, there is uncertainty in the cost associated to such actions, which is influenced by uncertain events that can be discovered over time. Thus, a common approach in such cases is to resort to *sampling* strategies, in order to obtain simple approximations of the probability distributions – in the form of sample realizations of the probabilistic costs. Due to the uncertainty arising in such problems, it is especially appealing to recur to solutions that adopt approaches to decision making under uncertainty, such as minimax, maximin, and regret-based methods.

As outlined in Table 3, more research effort is needed to solve DCOP models in which agents act in a combined uncertain and dynamic environment. This is by far the most realistic setting for teams of agents acting within a MAS. More generally, a coordination strategy that can adapt to the situation where the environment or network problem is evolving dynamically and rapidly and where several scenarios require different approaches to coordination, has not yet been studied.

Finally, *Quantified DCOPs* model adversarial agents, which are common in many MAS applications. In a Quantified DCOP, universally quantified variables can be considered as the choice of the nature of an adversary. Quantified DCOPs can thus formalize problems where a team of agents seeks to limit the effect of the adversarial agents, as well as problems associated to planning under uncertainty. Examples of relevant applications include *distributed surveillance planning problems*, where sensors in a network need to coordinate their surveillance areas to detect intruders, whose positions are unknown and who are trying to avoid detection. Thus, the sensors try to find a robust plan that can handle different intruding scenarios (see Section 8.5).

In terms of complexity, solving classical, Asymmetric, Multi-Objective, and Dynamic DCOPs optimally is NP-hard. In contrast, the Probabilistic DCOP extension to include the capability of handling uncertain events and the Quantified DCOP extension to model adversarial agents comes at a cost of increasing complexity. Both Probabilistic DCOPs and Quantified DCOPs are PSPACE-hard. Table 6 summarizes the complexity of the DCOP models discussed in this survey.

DCOP MODEL	COMPLEXITY
Classical	NP-hard
Asymmetric	NP-hard
Multi-objective	NP-hard
Dynamic	NP-hard
Probabilistic	PSPACE-hard
Quantified	PSPACE-hard

Table 6: Complexity of the DCOP Models

9.2 Algorithms and Theoretical Analysis

Despite the fact that classical DCOPs have reached a sufficient level of maturity from the algorithmic perspective, most of the other proposed formalisms fall short on both algorithmic and theoretical foundations. The proposed algorithms mostly extend classical DCOP algorithms and, therefore, they result in similar performance. Investigating strategies that are based on different backgrounds could help propel the evolution of this area.

9.2.1 INTER-DISCIPLINARY RESEARCH

Further investigations on relating DCOPs to the areas of game theory and decision theory are necessary. Similar to the work of Chapman, Rogers, and Jennings (2008), where the authors study relationship between DCOPs and potential games, analyzing the relationship of DCOPs to auction mechanisms could shed light on how to effectively address coordination and reasoning strategies in DCOP with partially cooperative agents. Another direction is to relate DCOP with machine learning techniques. For instance, as outlined by Kumar and Zilberstein (2011) and Ghosh, Kumar, and Varakantham (2015), one can use inference-based algorithms, such as expectation-maximization, and convex optimization machinery to develop efficient message-passing algorithms for solving large DCOPs. Additionally, merging insights from decision theory, such as handling partial observability, with the inherent DCOP ability of naturally exploiting problem structure could result in improved performance and/or refined models (Nair, Varakantham, Tambe, & Yokoo, 2005; Hoang et al., 2016, 2017).

9.2.2 ANYTIME MECHANISMS

Due to the complexity of the DCOP models, the study of incomplete approaches to solve large DCOPs whose agents may act in a dynamic and/or uncertain environment seems particularly suitable. Within the current incomplete methods proposed, a considerable effort has been employed in developing *anytime* algorithms (Zilberstein, 1996). In addition to the anytime mechanism, which constrains the problem resolution within a particular time requirement, *any-space* algorithms have been proposed to limit the amount of space needed to an agent during problem resolution (Petcu & Faltings, 2007a; Yeoh et al., 2009b, 2011; Gutierrez et al., 2011). Similar to these mechanisms, we see an urge for investigating algorithms that allow problem resolution within any particular network load restriction. For instance, in a congested network agents might need to exchange smaller messages, or to communicate less frequently and/or with a restricted set of neighbors. Situations like

these may arise in problems where a multitude of interconnected systems share the same medium, and thus limited bandwidth and interference may cause unsuitable delays. This context may be exacerbated with the advent of concepts such as the Internet of the Things, which expects to see million of connected devices, possibly sharing several mediums (Atzori, Iera, & Morabito, 2010; Chen, 2012; Miorandi, Sicari, De Pellegrini, & Chlamtac, 2012). Thus, orthogonal to the direction pursed by anytime and any-space algorithms, we envision the development of *any-communication* procedures.

9.2.3 COORDINATION AND COOPERATION STUDY

It has been observed that simple coordination strategies give good results in extremes (high or low) agent workload environments. However, at intermediate workload levels, such strategies lose their effectiveness, and more complex coordination strategies are necessary (Lesser, Decker, Wagner, Carver, Garvey, Horling, Neiman, Podorozhny, Prasad, Raja, et al., 2004; Zhang & Lesser, 2013). For example, Zhang and Lesser (2013) study agent learning, where coordination is driven by a dynamic decomposition of a DCOP, each solving smaller independent subproblems. Their proposal effectively produced near-optimal agent policies for learning and significantly reduced the amount of communication. These empirical observations suggest the existence of phase transition behaviors occurring at the level of agent coordination. A formal understanding of these phenomena, which looks at the environment, agents' local reasoning strategies, their assumptions on other agents states, as well as team coordination, could help researchers to understand the inherent complexity of coordination problems. Such a formal framework could be useful to build new and more efficient coordination strategies for a wide variety of multi-agent applications, perhaps resembling the way studies of phase transitions of NP-hard problems (Monasson, Zecchina, Kirkpatrick, Selman, & Troyansky, 1999) led to the understanding of problem complexity and creation of effective heuristics and search strategies.

9.3 DCOP Resolution Assumptions

Modeling many real-world problems as DCOPs often require each agent to solve large complex subproblems, each requiring many variables and constraints. A limitation of most DCOP algorithms is the assumption that each agent controls exactly one variable, and that all constraints are binary. Such assumptions simplify the algorithm organization and presentation. To operate under this assumption, reformulation techniques are commonly adopted to transform a general DCOP into one where each agent controls exclusively one variable. There are two commonly used reformulation techniques (Burke & Brown, 2006; Yokoo, 2001): *(i) Compilation* (or *complex variables*) where each agent creates a new *pseudo-variable* whose domain is the Cartesian product of the domains of all variables of the agent; and *(ii) Decomposition* (or *virtual agents*) where each agent creates a *pseudo-agent* for each of its variables. While both techniques are relatively simple, they can be inefficient. In compilation, the memory requirement for each agent grows exponentially with the number of variables that it controls. In decomposition, the DCOP algorithms will treat two pseudo-agents as independent entities, resulting in unnecessary computation and communication costs. A more realistic view is to allow each agent to solve its local subproblem (in a centralized fashion) since it is independent of the subproblems of other agents. The agent's subproblem resolution can then explore techniques from centralized reasoning, such as constraint optimization problems, linear programming, and graphical models. One can even exploit novel hardware platforms, such as

Graphical Processing Units (GPUs) to parallelize such solvers (Fioretto, Pontelli, Yeoh, & Dechter, 2018; Fioretto et al., 2016a; Bistaffa, Farinelli, & Bombieri, 2014; Fioretto, Campeotto, Fioretto, Yeoh, & Pontelli, 2014a).

Researchers have proposed re-designed versions of existing algorithms to be able to handle multi-variables agents problems (Davin & Modi, 2006; Khanna, Sattar, Hansen, & Stantic, 2009; Portway & Durfee, 2010). Additionally, recently, Grinshpoun (2015) and Fioretto, Yeoh, and Pontelli (2016b) have proposed general DCOP problem decompositions which are able to solve problems where agents control multiple variables.

9.4 Evaluation Metrics and Benchmarks

Another open question in this research area concerns the definition of a modeling language to model the different classes of DCOP problems and a systematic process to evaluate and compare the DCOP algorithms.

9.4.1 DCOP MODELING LANGUAGE

Despite the wide applicability of the DCOP model, unfortunately, there is no general language being used to formally specify a DCOP. While there are several DCOP simulators that include implementations of various DCOP algorithms using a common language specification (Sultanik, Modi, & Regli, 2007; Léauté, Ottens, & Szymanek, 2009; Wahbi, Ezzahir, Bessiere, & Bouyakhf, 2011), by and large, most stand-alone algorithms specify DCOPs in an ad-hoc manner. As a result, it is often inconvenient to experimentally compare different algorithms. More importantly, the great majority of such languages requires constraints values to be specified explicitly. Such requirement makes unpractical to convert problems which are naturally defined as mathematical optimization problems (such as Mixed Integer Programs) into an explicit form which specifies the utilities for each value combination of the variables in the scope of the problem constraints. The adoption of a common distributed constraint modeling language, which allows one to express constraints as standard algebraic or logic expressions, may be beneficial for developing standard benchmarks. Additionally, it would provide a tool for researchers outside the AI communities to model and test the applicability of new problems, extending the applicability of DCOPs to new areas. For instance, within the constraint programming community, the adoption of the MiniZinc language (Nethercote, Stuckey, Becket, Brand, Duck, & Tack, 2007) to model CSPs and COPs has gained wide traction, and it is becoming a modeling choice even outside the constraint programming community.

9.4.2 DCOP SIMULATORS

DCOP simulators are multi-agent software tools that simulate the execution of DCOP algorithms. They are a useful resource for DCOP researchers and practitioners for evaluating different models and algorithms. There are several MAS simulators which have been adopted by researchers to develop and compare DCOP algorithms:

- *DisChoco* (Wahbi et al., 2011) is a Java open-source framework for simulating DCOP algorithms where agents are executed asynchronously (each in a separate execution thread) and communication is implemented via messages passing. The simulator includes some problem generators and the possibility to take into account message loss, message corruption, and message delay. It supports models in which agents control multiple variables and problems with n-ary constraints. Additionally, it has the ability to be executed in a distributed environment.

- *Frodo* (Léauté et al., 2009) is a Java open-source framework for DCOP which implements several complete and incomplete classic DCOP algorithms, including DPOP and some of it variants, Max-Sum, MGM, and DSA. It can be executed either in a synchronous or in an asynchronous mode. It provides a Graphical User Interface to visualize the problem structure and algorithm execution, and it includes some problem generators for benchmarking. It supports problems with n-ary constraints and models in which agents need to control the assignment of several variables. Finally, in addition to simulating DCOP algorithms in a centralized environment, it provides support for executing agents in a distributed environment.
- *AgentZero* (Lutati, Gontmakher, Lando, Netzer, Meisels, & Grubshtein, 2014) is a Java open-source framework for various MAS applications. It provides distributed runtime environment, which can be used to simulate both synchronous and asynchronous algorithms. AgentZero offers several high-level APIs which allow to easily handle complex tasks, such as the construction of a pseudo-tree and the evaluation of assignments. Additionally, it provides some built-in performance measures and visualization tools that can be especially useful to debug distributed programs. It supports the notion of message delays, message corruptions, and message losses. While it does not support models in which agents control multiple variables, it can deal with n-ary constraints.

In addition to using DCOP simulators, researchers have developed and deployed DCOP algorithms using general MAS platforms. In particular, the *Java Agent DEvelopment Framework (JADE)* (Bellifemine, Caire, & Greenwood, 2007) is a multi-agent platform that is compliant with the FIPA specifications (FIPA, 2005). It includes a runtime environment in which the agents act, a set of APIs to define the agents' behaviors, and a suite of graphical tools that can be used to monitor the agents' activity. JADE agents' messages are specified by the ACL language, defined by the FIPA international standard for agent interoperability. JADE allows the deployment of agents on several environments, including mobile and wireless environments and fault-tolerant platforms.

Table 7 provides a summary of the simulators' features and describes whether the simulator *execution mode* supports synchronous or asynchronous algorithms; whether it provides *visualization* tools; the type of implementation of messages passing simulation (i.e., using pointers, deep copying the message structure into the receiving agent's message queue, etc.); if the framework provides support to simulate/handle message delays, message corruptions, and message losses; if it supports DCOP models in which agents control multiple variables and n-ary constraints; if the framework provides tools for debugging the algorithms; the programming language in which the simulator is implemented; whether it provides documentation; and whether it provides support for deployment. For a comprehensive review of multi-agent platforms, the interested reader is referred to the survey by Kravari and Bassiliades (2015).

The development of DCOP simulators have provided researchers with useful tools to develop new DCOP algorithms and to facilitate comparison of existing DCOP algorithms. However, all the DCOP simulators described above focus on the classical DCOP model, while Dynamic and Probabilistic DCOP simulators have received little attention. Thus, there is a need of developing high-fidelity simulators for Dynamic DCOPs, which simulate how the environment can vary over time, as well as for Probabilistic DCOPs, which simulate the stochastic nature of exogenous events.

PROPERTY	DISCHOCO	FRODO	AGENTZERO	JADE
Execution modes	Async	Sync/Async	Sync/Async	Synch/Async
Visualization	No	Yes	Yes	Yes
Message support	Pointers	Pointers	Deep copy	FIPA messages
Message delays/corruptions/losses	Yes	No	Yes	Yes
Multiple variables per agent support	Yes	Yes	No	N/A
N-ary constraints support	Yes	Yes	Yes	N/A
Debugging Tools	No	No	Yes	No
Programming Language	Java	Java	Java	Java
Documentation	No	Yes	Yes	Yes
Deployable	No	No	No	Yes

Table 7: Comparison of DCOP Simulators

9.4.3 EVALUATION AND METRICS

There are multiple metrics that can be used to measure the runtime of an algorithm, such as the number of *Non-Concurrent Constraint Checks (NCCCs)* (Meisels, Kaplansky, Razgon, & Zivan, 2002), the simulated runtime (Sultanik, Lass, & Regli, 2008), and the number of cycles (Lynch, 1996). However, there is no consensus on a standard metric. Such issues, combined with absence of a general DCOP language, make it inconvenient to experimentally compare different algorithms. In addition, new proposed algorithms are, in general, evaluated on arbitrary benchmarks, some inherited from the CSP literature, some other from approximations of real-world problems. To cope with these issues, it would be useful to develop a benchmark repository, perhaps by taking inspiration from the efforts made by the constraint programming community with CSPLib (CSPLib, 2000), a library of test problems for constraint solvers, or by the planning community with their international planning competitions (IPC, 1998).

10. Conclusions

DCOPs have emerged as a popular formalism for distributed reasoning and coordination in multi-agent systems. Due to its limitation to support complex, real-time, and uncertain environment, researchers have introduced several model extensions to handle dynamic and uncertain environments, as well as different levels of cooperation among the agents.

While DCOPs' theoretical foundation and algorithmic frameworks have matured significantly over the past decade, their applicability to realistic domains is lagging behind. This survey aims at linking the DCOP theoretical framework and solving strategies with a set of potential applications where its applicability is having or may have a significant impact.

This survey provided an analysis of the recent advances made by the AAMAS community within the DCOP framework and propose a categorization based on agent characteristics, environment properties, and type of teamwork adopted. Within the proposed classification, it *(i)* presented a review of the characteristics of the different algorithmic solutions; *(ii)* discussed a number of application domains that can be naturally modeled within each DCOP framework; and *(iii)* identified some potential directions for future work with regards to agent coordination, algorithm scalability, modeling languages, and evaluation criteria of DCOP models and algorithms.

Acknowledgments

References

Abounadi, J., Bertsekas, D., & Borkar, V. (2001). Learning algorithms for Markov decision processes with average cost. *SIAM Journal on Control and Optimization*, *40*(3), 681–698.

Agrawal, P., Kumar, A., & Varakantham, P. (2015). Near-optimal decentralized power supply restoration in smart grids. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1275–1283.

Aji, S. M., & McEliece, R. J. (2000). The generalized distributive law. *IEEE Transactions on Information Theory*, *46*(2), 325–343.

Akyildiz, I. F., Su, W., Sankarasubramaniam, Y., & Cayirci, E. (2002). Wireless sensor networks: a survey. *Computer Networks*, *38*(4), 393–422.

Amador, S., Okamoto, S., & Zivan, R. (2014). Dynamic multi-agent task allocation with spatial and temporal constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 1384–1390.

Amato, C., Chowdhary, G., Geramifard, A., Ure, N. K., & Kochenderfer, M. J. (2013). Decentralized control of partially observable Markov decision processes. In *Proceedings of the Annual Conference on Decision and Control (CDC)*, pp. 2398–2405.

Amigoni, F., Castelletti, A., & Giuliani, M. (2015). Modeling the management of water resources systems using Multi-Objective DCOPs. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 821–829.

Apt, K. (2003). *Principles of Constraint Programming*. Cambridge University Press.

Athanasiadis, D., Kockar, I., & McArthur, S. (2013). Distributed constraint optimisation for flexible network management. In *Proceedings of the IEEE International Conference on Innovative Smart Grid Technologies Europe (ISGT EUROPE)*, pp. 1–5.

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 (AAMAS)*, pp. 1047–1054.

Atzori, L., Iera, A., & Morabito, G. (2010). The internet of things: A survey. *Computer Networks*, *54*(15), 2787–2805.

Auer, P., Cesa-Bianchi, N., & Fischer, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine learning*, *47*(2-3), 235–256.

Baba, S., Iwasaki, A., Yokoo, M., Silaghi, M. C., Hirayama, K., & Matsui, T. (2010). Cooperative problem solving against adversary: Quantified distributed constraint satisfaction problem. In

Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 781–788.

Baba, S., Joe, Y., Iwasaki, A., & Yokoo, M. (2011). Real-time solving of quantified CSPs based on monte-carlo game tree search. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 655–661.

Bellifemine, F. L., Caire, G., & Greenwood, D. (2007). *Developing Multi-Agent Systems with JADE*, Vol. 7. John Wiley & Sons.

Benders, J. F. (1962). Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, *4*(1), 238–252.

Benedetti, M., Lallouet, A., & Vautard, J. (2008). Quantified constraint optimization. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*, pp. 463–477.

Bernstein, D. S., Givan, R., Immerman, N., & Zilberstein, S. (2002). The complexity of decentralized control of Markov decision processes. *Mathematics of Operations Research*, *27*(4), 819–840.

Bessiere, C., Brito, I., Gutierrez, P., & Meseguer, P. (2014). Global constraints in distributed constraint satisfaction and optimization. *Computer Journal*, *57*(6), 906–923.

Bessiere, C., Gutierrez, P., & Meseguer, P. (2012). Including soft global constraints in DCOPs. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*, pp. 175–190.

Billiau, G., Chang, C. F., & Ghose, A. (2012a). SBDO: A new robust approach to dynamic distributed constraint optimisation. In *Proceedings of the Principles and Practice of Multi-Agent Systems (PRIMA)*, pp. 11–26.

Billiau, G., Chang, C. F., Ghose, A., & Miller, A. A. (2012b). Using distributed agents for patient scheduling. In *Proceedings of the Principles and Practice of Multi-Agent Systems (PRIMA)*, pp. 551–560.

Billiau, G., Chang, C. F., Ghose, A., & Miller, A. (2012c). Support-based distributed optimisation: An approach to radiotherapy scheduling. In *Electronic Healthcare*, pp. 327–334.

Binmore, K. (1992). *Fun and Games, a Text on Game Theory*. DC Heath and Company.

Bistaffa, F., Farinelli, A., & Bombieri, N. (2014). Optimising memory management for belief propagation in junction trees using GPGPUs. In *Proceeding of the International Conference on Parallel and Distributed Systems (ICPADS)*, pp. 526–533.

Bowring, E., Tambe, M., & Yokoo, M. (2006). Multiply-constrained distributed constraint optimization. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1413–1420.

Brito, I., Meisels, A., Meseguer, P., & Zivan, R. (2009). Distributed constraint satisfaction with partially known constraints. *Constraints*, *14*(2), 199–234.

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

Burke, D., & Brown, K. (2006). Efficiently handling complex local problems in distributed constraint optimisation. In *Proceedings of the European Conference on Artificial Intelligence (ECAI)*, pp. 701–702.

Carpenter, C. J., Dugan, C. J., Kopena, J. B., Lass, R. N., Naik, G., Nguyen, D. N., Sultanik, E., Modi, P. J., & Regli, W. C. (2007). Intelligent systems demonstration: disaster evacuation support. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 1964–1965.

Cerquides, J., Picard, G., & Rodríguez-Aguilar, J. A. (2015). Designing a marketplace for the trading and distribution of energy in the smart grid. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1285–1293.

Chachra, S., & Marefat, M. (2006). Distributed algorithms for sleep scheduling in wireless sensor networks. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3101–3107.

Chapman, A., Rogers, A., & Jennings, N. (2008). A parameterisation of algorithms for distributed constraint optimisation via potential games. In *International Workshop on Distributed Constraint Reasoning (DCR)*, pp. 99–113.

Chechetka, A., & Sycara, K. (2006). No-commitment branch and bound search for distributed constraint optimization. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1427–1429.

Chen, Y.-K. (2012). Challenges and opportunities of internet of things. In *Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC)*, pp. 383–388.

Chen, Z., Deng, Y., & Wu, T. (2017). An iterative refined max-sum_ad algorithm via single-side value propagation and local search. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 195–202.

Chli, M., & Winsper, M. (2015). Using the Max-Sum algorithm for supply chain emergence in dynamic multiunit environments. *IEEE Transactions of Systems, Man, and Cybernetics: Systems*, *45*(3), 422–435.

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. The PAAMS Collection*, pp. 50–61.

Choxi, H., & Modi, P. (2007). A distributed constraint optimization approach to wireless network optimization. In *Proceedings of the AAAI-07 Workshop on Configuration*, pp. 1–8.

CSPLib (2000). CSPLib: A problem library for constraints. `http://www.csplib.org/`.

Davidson, E., Dolan, M., McArthur, S., & Ault, G. (2009). The use of constraint programming for the autonomous management of power flows. In *Proceedings of the International Conference on Intelligent System Applications to Power Systems (ISAP)*, pp. 1–7.

Davin, J., & Modi, P. (2006). Hierarchical variable ordering for multiagent agreement problems. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1433–1435.

De Oliveira, D., Bazzan, A. L., & Lesser, V. (2005). Using cooperative mediation to coordinate traffic lights: a case study. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 463–470.

Delle Fave, F. M., Stranders, R., Rogers, A., & Jennings, N. R. (2011). Bounded decentralised coordination over multiple objectives. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 371–378.

Dijkstra, E. W. (1974). Self-stabilization in spite of distributed control. *Communication of the ACM*, *17*(11), 643–644.

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 (AAMAS)*, pp. 639–646.

Fioretto, F., Campeotto, F., Fioretto, L. D. R., Yeoh, W., & Pontelli, E. (2014a). GD-GIBBS: a GPU-based sampling algorithm for solving distributed constraint optimization problems. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1339–1340.

Fioretto, F., Le, T., Yeoh, W., Pontelli, E., & Son, T. C. (2014b). Improving DPOP with branch consistency for solving distributed constraint optimization problems. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*, pp. 307–323.

Fioretto, F., Pontelli, E., Yeoh, W., & Dechter, R. (2018). Accelerating exact and approximate inference for distributed) discrete optimization with GPUs. *Constraints*, *23*(1), 1–43.

Fioretto, F., Yeoh, W., & Pontelli, E. (2016a). A dynamic programming-based MCMC framework for solving DCOPs with GPUs. In *Proceedings of Principles and Practice of Constraint Programming (CP)*, pp. 813–831.

Fioretto, F., Yeoh, W., & Pontelli, E. (2016b). Multi-variable agents decomposition for DCOPs. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 2480–2486.

Fioretto, F., Yeoh, W., & Pontelli, E. (2017a). A multiagent system approach to scheduling devices in smart homes. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 981–989.

Fioretto, F., Yeoh, W., Pontelli, E., Ma, Y., & Ranade, S. (2017b). A DCOP approach to the economic dispatch with demand response. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 999–1007.

FIPA (2005). The Foundation for Intelligent Physical Agents. http://www.fipa.org/.

Garrido, L., & Sycara, K. (1996). Multi-agent meeting scheduling: Preliminary experimental results. In *Proceedings of the International Conference on Multiagent Systems*, pp. 95–102.

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

Geman, S., & Geman, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *6*(6), 721–741.

Gershman, A., Meisels, A., & Zivan, R. (2009). Asynchronous Forward-Bounding for distributed COPs. *Journal of Artificial Intelligence Research*, *34*, 61–88.

Ghosh, S., Kumar, A., & Varakantham, P. (2015). Probabilistic inference based message-passing for resource constrained DCOPs. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 411–417.

Giffler, B., & Thompson, G. L. (1960). Algorithms for solving production-scheduling problems. *Operations research*, *8*(4), 487–503.

Giuliani, M., Castelletti, A., Amigoni, F., & Cai, X. (2014). Multiagent systems and distributed constraint reasoning for regulatory mechanism design in water management. *Journal of Water Resources Planning and Management*, *41*(4).

Glaßer, C., Reitwießner, C., Schmitz, H., & Witek, M. (2010). Approximability and hardness in multi-objective optimization. In *Proceedings of the Conference on Computability in Europe (CiE)*, pp. 180–189.

Golomb, S. W., & Baumert, L. D. (1965). Backtrack programming. *Journal of the ACM*, *12*(4), 516–524.

Greenstadt, R., Grosz, B., & Smith, M. (2007). SSDPOP: Improving the privacy of DCOP with secret sharing. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1098–1100.

Grinshpoun, T. (2015). Clustering variables by their agents. In *Proceedings of the International Joint Conferences on Web Intelligence and Intelligent Agent Technologies (WI-IAT)*, Vol. 2, pp. 250–256.

Grinshpoun, T., Grubshtein, A., Zivan, R., Netzer, A., & Meisels, A. (2013). Asymmetric distributed constraint optimization problems. *Journal of Artificial Intelligence Research*, *47*, 613–647.

Grinshpoun, T., & Meisels, A. (2008). Completeness and performance of the APO algorithm. *Journal of Artificial Intelligence Research*, *33*, 223–258.

Gupta, S., Jain, P., Yeoh, W., Ranade, S., & Pontelli, E. (2013a). Solving customer-driven microgrid optimization problems as DCOPs. In *International Workshop on Distributed Constraint Reasoning (DCR)*, pp. 45–59.

Gupta, S., Yeoh, W., Pontelli, E., Jain, P., & Ranade, S. (2013b). Modeling microgrid islanding problems as DCOPs. In *Proceedings of the North American Power Symposium (NAPS)*, pp. 1–6.

Gutierrez, P., Lee, J., Lei, K. M., Mak, T., & Meseguer, P. (2013). Maintaining soft arc consistencies in BnB-ADOPT$^+$ during search. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*, pp. 365–380.

Gutierrez, P., & Meseguer, P. (2012a). Improving BnB-ADOPT$^+$-AC. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 273–280.

Gutierrez, P., & Meseguer, P. (2012b). Removing redundant messages in n-ary BnB-ADOPT. *Journal of Artificial Intelligence Research*, *45*, 287–304.

Gutierrez, P., Meseguer, P., & Yeoh, W. (2011). Generalizing ADOPT and BnB-ADOPT. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 554–559.

Hannebauer, M., & Müller, S. (2001). Distributed constraint optimization for medical appointment scheduling. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 139–140.

Harvey, P., Chang, C. F., & Ghose, A. (2007). Support-based distributed search: a new approach for multiagent constraint processing. In *Argumentation in Multi-Agent Systems*, pp. 91–106. Springer Berlin Heidelberg.

Hirayama, K., & Yokoo, M. (1997). Distributed partial constraint satisfaction problem. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*, pp. 222–236.

Hoang, K. D., Fioretto, F., Hou, P., Yokoo, M., Yeoh, W., & Zivan, R. (2016). Proactive dynamic distributed constraint optimization. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 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 Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 212–220.

Hollos, D., Karl, H., & Wolisz, A. (2004). Regionalizing global optimization algorithms to improve the operation of large ad hoc networks. In *Proceedings of the Wireless Communications and Networking Conference (WCNC)*, pp. 819–824.

IPC (1998). International Planning Competition. `http://www.icaps-conference.org/index.php/Main/Competitions`.

Jain, M., Taylor, M. E., Tambe, M., & Yokoo, M. (2009). DCOPs meet the real world: Exploring unknown reward matrices with applications to mobile sensor networks. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 181–186.

Jain, P., Ranade, S. J., Gupta, S., & Pontelli, E. (2012). Optimum operation of a customer-driven microgrid: A comprehensive approach. In *Proceedings of the IEEE International Conference on Power Electronics, Drives and Energy Systems (PEDES)*, pp. 1–6.

Jennings, N. R., & Jackson, A. (1995). Agent-based meeting scheduling: A design and implementation. *Electronics letters*, *31*(5), 350–352.

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 (DASC)*, 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 (AAMAS)*, pp. 599–606.

Kearns, M., Littman, M. L., & Singh, S. (2001). Graphical models for game theory. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 253–260.

Khanna, S., Sattar, A., Hansen, D., & Stantic, B. (2009). An efficient algorithm for solving dynamic complex DCOP problems. In *Proceedings of the International Joint Conferences on Web Intelligence and Intelligent Agent Technologies (WI-IAT)*, pp. 339–346.

Kiekintveld, C., Yin, Z., Kumar, A., & Tambe, M. (2010). Asynchronous algorithms for approximate distributed constraint optimization with quality bounds. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 133–140.

Kinoshita, K., Iizuka, K., & Iizuka, Y. (2013). Effective disaster evacuation by solving the distributed constraint optimization problem. In *Proceedings of the International Conference on Advanced Applied Informatics (IIAIAAI)*, pp. 399–400.

Kluegel, W., Iqbal, M. A., Fioretto, F., Yeoh, W., & Pontelli, E. (2017). A realistic dataset for the smart home device scheduling problem for DCOPs. In Sukthankar, G., & Rodriguez-Aguilar, J. A. (Eds.), *Autonomous Agents and Multiagent Systems: AAMAS 2017 Workshops, Visionary Papers, São Paulo, Brazil, May 8-12, 2017, Revised Selected Papers*, pp. 125–142. Springer International Publishing.

Kopena, J. B., Sultanik, E. A., Lass, R. N., Nguyen, D. N., Dugan, C. J., Modi, P. J., & Regli, W. C. (2008). Distributed coordination of first responders. *Internet Computing*, *12*(1), 45–47.

Kraus, S. (1997). Negotiation and cooperation in multi-agent environments. *Artificial Intelligence*, *94*(1), 79–97.

Kravari, K., & Bassiliades, N. (2015). A survey of agent platforms. *Journal of Artificial Societies and Social Simulation*, *18*(1), 11.

Kschischang, F. R., Frey, B. J., & Loeliger, H.-A. (2001). Factor graphs and the sum-product algorithm. *IEEE Transactions on Information Theory*, *47*(2), 498–519.

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 (AAMAS)*, pp. 923–930.

Kumar, A., Petcu, A., & Faltings, B. (2008). H-DPOP: Using hard constraints for search space pruning in DCOP. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 325–330.

Kumar, A., & Zilberstein, S. (2011). Message-passing algorithms for quadratic programming formulations of map estimation. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 428–435.

Lallouet, A., Lee, J. H. M., Mak, T. W. K., & Yip, J. (2015). Ultra-weak solutions and consistency enforcement in minimax weighted constraint satisfaction. *Constraints*, *20*(2), 109–154.

Larrosa, J. (2002). Node and arc consistency in weighted CSP. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 48–53.

Lass, R. N., Kopena, J. B., Sultanik, E. A., Nguyen, D. N., Dugan, C. P., Modi, P. J., & Regli, W. C. (2008a). Coordination of first responders under communication and resource constraints. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1409–1412.

Lass, R. N., Regli, W. C., Kaplan, A., Mitkus, M., & Sim, J. J. (2008b). Facilitating communication for first responders using dynamic distributed constraint optimization. In *Proceedings of the Symposium on Technologies for Homeland Security (IEEE HST)*, pp. 316–320.

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 (AAMAS), pp. 606–614.

Le, T., Son, T. C., Pontelli, E., & Yeoh, W. (2015). Solving distributed constraint optimization problems with logic programming. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 1174–1181.

Léauté, T., & Faltings, B. (2009). E [DPOP]: Distributed constraint optimization under stochastic uncertainty using collaborative sampling. In *International Workshop on Distributed Constraint Reasoning (DCR)*, pp. 87–101.

Léauté, T., & Faltings, B. (2011). Distributed constraint optimization under stochastic uncertainty. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 68–73.

Léauté, T., Ottens, B., & Szymanek, R. (2009). FRODO 2.0: An open-source framework for distributed constraint optimization. In *International Workshop on Distributed Constraint Reasoning (DCR)*, pp. 160–164.

Lesser, V., Decker, K., Wagner, T., Carver, N., Garvey, A., Horling, B., Neiman, D., Podorozhny, R., Prasad, M. N., Raja, A., et al. (2004). Evolution of the GPGP/TAEMS domain-independent coordination framework. *Journal of Autonomous Agents and Multi-Agent Systems*, *9*(1-2), 87–143.

Levit, V., Grinshpoun, T., Meisels, A., & Bazzan, A. L. (2013). Taxation search in boolean games. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 183–190.

Li, S., Negenborn, R. R., & Lodewijks, G. (2016). Distributed constraint optimization for addressing vessel rotation planning problems. *Engineering Applications of Artificial Intelligence*, *48*, 159–172.

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.

Lorenzi, F., dos Santos, F., Ferreira Jr, P. R., & Bazzan, A. L. (2008). Optimizing preferences within groups: A case study on travel recommendation. In *Proceedings of the Brazilian Symposium on Artificial Intelligence (SBIA)*, pp. 103–112.

Lutati, B., Gontmakher, I., Lando, M., Netzer, A., Meisels, A., & Grubshtein, A. (2014). AgentZero: A framework for simulating and evaluating multi-agent algorithms. In *Agent-Oriented Software Engineering*, pp. 309–327. Springer Berlin Heidelberg.

Lynch, N. A. (1996). *Distributed Algorithms*. Morgan Kaufmann.

Macarthur, K., Farinelli, A., Ramchurn, S., & Jennings, N. (2010). Efficient, superstabilizing decentralised optimisation for dynamic task allocation environments. In *International Workshop on Optimization In Multi-Agent Systems (OPTMAS)*, pp. 25–32.

Macarthur, K., Farinelli, A., Ramchurn, S., & Jennings, N. (2011). A distributed anytime algorithm for dynamic task allocation in multi-agent systems. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 701–706.

Mackworth, A. K., & Freuder, E. C. (1985). The complexity of some polynomial network consistency algorithms for constraint satisfaction problems. *Artificial intelligence*, *25*(1), 65–74.

Mahadevan, S. (1996). Average reward reinforcement learning: Foundations, algorithms, and empirical results. *Machine Learning*, *22*(1–3), 159–195.

Maheswaran, R., Pearce, J., & Tambe, M. (2004a). Distributed algorithms for DCOP: A graphical game-based approach. In *Proceedings of the International Conference on Parallel and Distributed Computing Systems (PDCS)*, pp. 432–439.

Maheswaran, R. T., Tambe, M., Bowring, E., Pearce, J. P., & Varakantham, P. (2004b). Taking DCOP to the real world: Efficient complete solutions for distributed multi-event scheduling. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 310–317.

Mailler, R., & Lesser, V. (2004). Solving distributed constraint optimization problems using cooperative mediation. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 438–445.

Marler, R. T., & Arora, J. S. (2004). Survey of multi-objective optimization methods for engineering. *Structural and Multidisciplinary Optimization*, *26*(6), 369–395.

Matsui, T., Matsuo, H., Silaghi, M. C., Hirayama, K., Yokoo, M., & Baba, S. (2010). A quantified distributed constraint optimization problem. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1023–1030.

Matsui, T., & Matsuo, H. (2008). A formalization for distributed cooperative sensor resource allocation. In *Agent and Multi-Agent Systems: Technologies and Applications*, pp. 292–301.

Matsui, T., & Matsuo, H. (2011). A distributed cooperative model for resource supply networks. In *Proceedings of the International MultiConference of Engineers and Computer Scientists*, pp. 7–14.

Matsui, T., Matsuo, H., Silaghi, M., Hirayama, K., & Yokoo, M. (2008). Resource constrained distributed constraint optimization with virtual variables. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 120–125.

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 (PRIMA)*, pp. 137–152.

Matsui, T., Silaghi, M., Hirayama, K., Yokoo, M., & Matsuo, H. (2014). Leximin multiple objective optimization for preferences of agents. In *Proceedings of the Principles and Practice of Multi-Agent Systems (PRIMA)*, pp. 423–438.

Medi, A., Okimoto, T., & Inoue, K. (2014). A two-phase complete algorithm for multi-objective distributed constraint optimization. *Journal of Advanced Computational Intelligence and Intelligent Informatics*, *18*(4), 573–580.

Meisels, A. (2008). *Distributed Search by Constrained Agents: Algorithms, Performance, Communication*. Springer Science & Business Media.

Meisels, A., Kaplansky, E., Razgon, I., & Zivan, R. (2002). Comparing performance of distributed constraints processing algorithms. In *International Workshop on Distributed Constraint Reasoning (DCR)*, pp. 86–93.

Mejias, B., & Roy, P. V. (2010). From mini-clouds to cloud computing. In *International Conference on Self-Adaptive and Self-Organizing Systems Workshop (SASOW)*, pp. 234–238.

Miettinen, K. (1999). *Nonlinear Multiobjective Optimization*, Vol. 12. Springer Berlin Heidelberg.

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 (AAMAS)*, pp. 281–288.

Minton, S., Johnston, M. D., Philips, A. B., & Laird, P. (1992). Minimizing conflicts: a heuristic repair method for constraint satisfaction and scheduling problems. *Artificial Intelligence*, *58*(1), 161–205.

Miorandi, D., Sicari, S., De Pellegrini, F., & Chlamtac, I. (2012). Internet of things: Vision, applications and research challenges. *Ad Hoc Networks*, *10*(7), 1497–1516.

Mir, U., Merghem-Boulahia, L., & Gaïti, D. (2010). A cooperative multiagent based spectrum sharing. In *Advanced International Conference on Telecommunications (AICT)*, pp. 124–130.

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.

Molisch, A. F. (2012). *Wireless Communications*, Vol. 34. John Wiley & Sons.

Monasson, R., Zecchina, R., Kirkpatrick, S., Selman, B., & Troyansky, L. (1999). Determining computational complexity from characteristic phase transitions. *Nature*, *400*(6740), 133–137.

Monteiro, T. L., Pellenz, M. E., Penna, M. C., Enembreck, F., Souza, R. D., & Pujolle, G. (2012a). Channel allocation algorithms for WLANs using distributed optimization. *AEU-International Journal of Electronics and Communications*, *66*(6), 480–490.

Monteiro, T. L., Pujolle, G., Pellenz, M. E., Penna, M. C., Enembreck, F., & Demo Souza, R. (2012b). A multi-agent approach to optimal channel assignment in WLANs. In *Proceedings of the Wireless Communications and Networking Conference (WCNC)*, pp. 2637–2642.

Moreno-Vozmediano, R., Montero, R. S., & Llorente, I. M. (2013). Key challenges in cloud computing: Enabling the future internet of services. *Internet Computing*, *17*(4), 18–25.

Nair, R., Varakantham, P., Tambe, M., & Yokoo, M. (2005). Networked distributed POMDPs: A synergy of distributed constraint optimization and POMDPs. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 1758–1760.

Nethercote, N., Stuckey, P. J., Becket, R., Brand, S., Duck, G. J., & Tack, G. (2007). Minizinc: Towards a standard CP modelling language. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*, pp. 529–543.

Netzer, A., Grubshtein, A., & Meisels, A. (2012). Concurrent forward bounding for distributed constraint optimization problems. *Artificial Intelligence*, *193*, 186–216.

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 (AAMAS)*, 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 (AAMAS)*, 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 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1341–1342.

Noriega, P., & Sierra, C. (1999). Auctions and multi-agent systems. In *Intelligent Information Agents*, pp. 153–175. Springer.

Okimoto, T., Clement, M., & Inoue, K. (2013). AOF-based algorithm for dynamic Multi-Objective distributed constraint optimization. In *Multi-disciplinary Trends in Artificial Intelligence (MI-WAI)*, pp. 175–186.

Okimoto, T., Schwind, N., Clement, M., & Inoue, K. (2014). Lp-Norm based algorithm for multi-objective distributed constraint optimization (Extended Abstract). In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1427–1428.

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 (PRIMA)*, pp. 245–260.

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.

Paquete, L., Chiarandini, M., & Sttzle, T. (2004). Pareto local optimum sets in the biobjective traveling salesman problem: An experimental study. In *Metaheuristics for Multiobjective Optimisation*, Vol. 535, pp. 177–199.

Parsons, S., & Wooldridge, M. (2002). Game theory and decision theory in multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 5(3), 243–254.

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 (IJCAI)*, pp. 1446–1451.

Pecora, F., Modi, P., & Scerri, P. (2006). Reasoning about and dynamically posting n-ary constraints in adopt. In *International Workshop on Distributed Constraint Reasoning (DCR)*.

Penya-Alba, T., Cerquides, J., Rodriguez-Aguilar, J. A., & Vinyals, M. (2012). Scalable decentralized supply chain formation through binarized belief propagation. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1275–1276.

Penya-Alba, T., Vinyals, M., Cerquides, J., & Rodriguez-Aguilar, J. A. (2014). Exploiting Max-Sum for the decentralized assembly of high-valued supply chains. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 373–380.

Penya-Alba, T., Vinyals, M., Cerquides, J., & Rodriguez-Aguilar, J. A. (2012). A scalable message-passing algorithm for supply chain formation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 1436–1442.

Peri, O., & Meisels, A. (2013). Synchronizing for performance-DCOP algorithms. In *Proceedings of the International Conference on Agents and Artificial Intelligence (ICAART)*, pp. 5–14.

Petcu, A., & Faltings, B. (2005a). Approximations in distributed optimization. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*, pp. 802–806.

Petcu, A., & Faltings, B. (2005b). A scalable method for multiagent constraint optimization. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 1413–1420.

Petcu, A., & Faltings, B. (2005c). Superstabilizing, fault-containing distributed combinatorial optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 449–454.

Petcu, A., & Faltings, B. (2006). ODPOP: An algorithm for open/distributed constraint optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 703–708.

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 (IJCAI)*, 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 (IAT)*, pp. 321–327.

Petcu, A., Faltings, B., & Mailler, R. (2007). PC-DPOP: A new partial centralization algorithm for distributed optimization. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 167–172.

Pham, T., Tawfik, A., & Taylor, M. (2013). A simple, naive agent-based model for the optimization of a system of traffic lights: Insights from an exploratory experiment. In *Proceedings of the Conference on Agent-based Modeling in Transportation Planning and Operations*, pp. 1–21.

Portway, C., & Durfee, E. H. (2010). The multi variable multi constrained distributed constraint optimization framework. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1385–1386.

Pujol-Gonzalez, M., Cerquides, J., Farinelli, A., Meseguer, P., & Rodriguez-Aguilar, J. A. (2015). Efficient inter-team task allocation in RoboCup rescue. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 413–421.

Raiffa, H. (1968). *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Addison-Wesley.

Ramchurn, S. D., Farinelli, A., Macarthur, K. S., & Jennings, N. R. (2010). Decentralized coordination in RoboCup rescue. *Computer Journal*, *53*(9), 1447–1461.

Ramchurn, S. D., Vytelingum, P., Rogers, A., & Jennings, N. R. (2012). Putting the 'smarts' into the smart grid: A grand challenge for artificial intelligence. *Communications of the ACM*, *55*(4), 86–97.

Regan, K., & Boutilier, C. (2010). Robust policy computation in reward-uncertain MDPs using nondominated policies. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 1127–1133.

Rogers, A., Farinelli, A., Stranders, R., & Jennings, N. (2011). Bounded approximate decentralised coordination via the Max-Sum algorithm. *Artificial Intelligence*, *175*(2), 730–759.

Rollon, E., & Larrosa, J. (2012). Improved Bounded Max-Sum for distributed constraint optimization. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*, pp. 624–632.

Rossi, F., Beek, P. v., & Walsh, T. (2006). *Handbook of Constraint Programming (Foundations of Artificial Intelligence)*. Elsevier Science Inc.

Rust, P., Picard, G., & Ramparany, F. (2016). Using message-passing DCOP algorithms to solve energy-efficient smart environment configuration problems. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 468–474.

Salukvad, M. E. (1971). Optimization of vector functionals. I. programming of optimal trajectories. *Automation and remote Control*, *32*(8), 5–15.

Schwartz, A. (1993). A reinforcement learning method for maximizing undiscounted rewards. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 298–305.

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

Shapiro, L. G., & Haralick, R. M. (1981). Structural descriptions and inexact matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *3*(5), 504–519.

Solomon, M. M. (1987). Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research*, *35*(2), 254–265.

Steinbauer, G., & Kleiner, A. (2012). Towards CSP-based mission dispatching in C2/C4I systems. In *IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, pp. 1–6.

Stranders, R., Delle Fave, F. M., Rogers, A., & Jennings, N. (2011). U-GDL: A decentralised algorithm for DCOPs with uncertainty. Tech. rep., University of Southampton, Department of Electronics and Computer Science.

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 (AAMAS)*, pp. 601–608.

Stranders, R., Tran-Thanh, L., Fave, F. M. D., Rogers, A., & Jennings, N. R. (2012). DCOPs and bandits: Exploration and exploitation in decentralised coordination. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 289–296.

Sultanik, E., Modi, P. J., & Regli, W. C. (2007). On modeling multiagent task scheduling as a distributed constraint optimization problem. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 1531–1536.

Sultanik, E. A., Lass, R. N., & Regli, W. C. (2008). DCOPolis: A framework for simulating and deploying distributed constraint reasoning algorithms. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1667–1668. International Foundation for Autonomous Agents and Multiagent Systems.

Tabakhi, A. M., Le, T., Fioretto, F., & Yeoh, W. (2017). Preference elicitation for DCOPs. In *Proceedings of Principles and Practice of Constraint Programming (CP)*, pp. 278–296.

Tassa, T., Zivan, R., & Grinshpoun, T. (2016). Preserving privacy in region optimal DCOP algorithms. In *IJCAI*, pp. 496–502.

Taylor, M. E., Jain, M., Jin, Y., Yokoo, M., & Tambe, M. (2010). When should there be a me in team?: Distributed multi-agent optimization under uncertainty. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 109–116.

Taylor, M. E., Jain, M., Tandon, P., Yokoo, M., & Tambe, M. (2011). Distributed on-line multi-agent optimization under uncertainty: Balancing exploration and exploitation. *Advances in Complex Systems*, *14*(03), 471–528.

Van Hentenryck, P., & Michel, L. (2009). *Constraint-based local search*. The MIT Press.

Van Katwijk, R., De Schutter, B., & Hellendoorn, J. (2009). Multi-agent control of traffic networks: Algorithm and case study. In *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1–6.

Verfaillie, G., & Jussien, N. (2005). Constraint solving in uncertain and dynamic environments: A survey. *Constraints*, *10*(3), 253–281.

Vermorel, J., & Mohri, M. (2005). Multi-armed bandit algorithms and empirical evaluation. In *Proceedings of the European Conference on Machine Learning (ECML)*, pp. 437–448. Springer.

Vinyals, M., Shieh, E., Cerquides, J., Rodriguez-Aguilar, J., Yin, Z., Tambe, M., & Bowring, E. (2011). Quality guarantees for region optimal DCOP algorithms. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 133–140.

Wack, M., Okimoto, T., Clement, M., & Inoue, K. (2014). Local search based approximate algorithm for Multi-Objective DCOPs. In *Proceedings of the Principles and Practice of Multi-Agent Systems (PRIMA)*, pp. 390–406.

Wahbi, M., Ezzahir, R., Bessiere, C., & Bouyakhf, E.-H. (2011). Dischoco 2: A platform for distributed constraint reasoning. *International Workshop on Distributed Constraint Reasoning (DCR)*, *11*, 112–121.

Walsh, W. E., & Wellman, M. P. (2000). Modeling supply chain formation in multiagent systems. In *Agent Mediated Electronic Commerce II*, pp. 94–101. Springer.

Wang, Y., Sycara, K., & Scerri, P. (2011). Towards an understanding of the value of cooperation in uncertain world. In *Proceedings of the International Joint Conferences on Web Intelligence and Intelligent Agent Technologies (WI-IAT)*, Vol. 2, pp. 212–215.

Winsper, M., & Chli, M. (2013). Decentralized supply chain formation using Max-Sum loopy belief propagation. *Computational Intelligence*, *29*(2), 281–309.

Wood, A. J., & Wollenberg, B. F. (2012). *Power Generation, Operation, and Control*. John Wiley & Sons.

Wooldridge, M. (2009). *An introduction to multiagent systems*. John Wiley & Sons.

Wu, F., & Jennings, N. R. (2014). Regret-based multi-agent coordination with uncertain task rewards. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 1492–1499.

Xie, J., Howitt, I., & Raja, A. (2007). Cognitive radio resource management using multi-agent systems. In *Proceedings of the Consumer Communications and Networking Conference (CCNC)*, pp. 1123–1127.

Yan, Z., Lee, J.-H., Shen, S., & Qiao, C. (2013). Novel branching-router-based multicast routing protocol with mobility support. *IEEE Transactions on Parallel and Distributed Systems*, *24*(10), 2060–2068.

Yedidsion, H., & Zivan, R. (2014). Applying DCOP_MST to a team of mobile robots with directional sensing abilities. In *International Joint Workshop on Optimization In Multi-Agent Systems and Distributed Constraint Reasoning (OPTMAS-DCR)*.

Yedidsion, H., Zivan, R., & Farinelli, A. (2014). Explorative Max-Sum for teams of mobile sensing agents. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 549–556.

Yeoh, W. (2010). *Speeding Up Distributed Constraint Optimization Search Algorithms*. Ph.D. thesis, University of Southern California, Los Angeles (United States).

Yeoh, W., Felner, A., & Koenig, S. (2010). BnB-ADOPT: An asynchronous branch-and-bound DCOP algorithm. *Journal of Artificial Intelligence Research*, *38*, 85–133.

Yeoh, W., Sun, X., & Koenig, S. (2009a). Trading off solution quality for faster computation in DCOP search algorithms. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 354–360.

Yeoh, W., Varakantham, P., & Koenig, S. (2009b). Caching schemes for DCOP search algorithms. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 609–616.

Yeoh, W., Varakantham, P., Sun, X., & Koenig, S. (2011). Incremental DCOP search algorithms for solving dynamic DCOPs (Extended Abstract). In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1069–1070.

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

Yokoo, M. (Ed.). (2001). *Distributed constraint satisfaction: Foundation of cooperation in Multi-agent Systems*. Springer Berlin Heidelberg.

Yokoo, M., Durfee, E. H., Ishida, T., & Kuwabara, K. (1998). The distributed constraint satisfaction problem: Formalization and algorithms. *Knowledge and Data Engineering, IEEE Transactions on*, *10*(5), 673–685.

Yokoo, M., & Hirayama, K. (2000). Algorithms for distributed constraint satisfaction: A review. *Autonomous Agents and Multi-Agent Systems*, *3*(2), 185–207.

Zhang, C., & Lesser, V. (2013). Coordinating multi-agent reinforcement learning with limited communication. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1101–1108.

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.

Zilberstein, S. (1996). Using anytime algorithms in intelligent systems. *AI Magazine*, *17*(3), 73.

Zivan, R., Glinton, R., & Sycara, K. (2009). Distributed constraint optimization for large teams of mobile sensing agents. In *Proceedings of the International Joint Conferences on Web Intelligence and Intelligent Agent Technologies (WI-IAT)*, pp. 347–354.

Zivan, R., & Meisels, A. (2006). Dynamic ordering for asynchronous backtracking on DisCSPs. *Constraints*, *11*(2-3), 179–197.

Zivan, R., Okamoto, S., & Peled, H. (2014). Explorative anytime local search for distributed constraint optimization. *Artificial Intelligence*, *212*, 1–26.

Zivan, R., Parash, T., Cohen, L., Peled, H., & Okamoto, S. (2017). Balancing exploration and exploitation in incomplete Min/Max-Sum inference for distributed constraint optimization. *Journal of Autonomous Agents and Multi-Agent Systems*, 1–43.

Zivan, R., & Peled, H. (2012). Max/Min-Sum distributed constraint optimization through value propagation on an alternating DAG. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 265–272.