

# *Solving distributed constraint optimization problems using logic programming\*†*

TIEP LE, TRAN CAO SON, ENRICO PONTELLI and WILLIAM YEOH

*Computer Science Department, New Mexico State University  
Las Cruces, NM 88001, USA*

(e-mails: tile@cs.nmsu.edu, tson@cs.nmsu.edu, epontell@cs.nmsu.edu,  
wyeoh@cs.nmsu.edu)

*submitted 5 April 2016; revised 10 May 2017; accepted 29 May 2017*

---

## Abstract

This paper explores the use of *Answer Set Programming (ASP)* in solving *Distributed Constraint Optimization Problems (DCOPs)*. The paper provides the following novel contributions: (1) it shows how one can formulate DCOPs as logic programs; (2) it introduces ASP-DPOP, the first DCOP algorithm that is based on logic programming; (3) it experimentally shows that ASP-DPOP can be up to two orders of magnitude faster than DPOP (its imperative programming counterpart) as well as solve some problems that DPOP fails to solve, due to memory limitations; and (4) it demonstrates the applicability of ASP in a wide array of multi-agent problems currently modeled as DCOPs.

**KEYWORDS:** DCOP, DPOP, logic programming, ASP

---

## 1 Introduction

Distributed constraint optimization problems (DCOPs) are optimization problems where agents need to coordinate the assignment of values to their “local” variables to maximize the overall sum of resulting constraint utilities (Mailler and Lesser 2004; Modi *et al.* 2005; Petcu and Faltings 2005a; Yeoh and Yokoo 2012). The process is subject to limitations on the communication capabilities of the agents; in particular, each agent can only exchange information with neighboring agents within a given topology. DCOPs are well suited for modeling multi-agent coordination and resource allocation problems, where the primary interactions are between local subsets of agents. Researchers have used DCOPs to model various problems, such as the distributed scheduling of meetings (Maheswaran *et al.* 2004; Zivan *et al.* 2014), distributed allocation of targets to sensors in a network (Farinelli *et al.* 2008), distributed allocation of resources in disaster evacuation scenarios (Lass *et al.* 2008),

\* This article extends our previous conference paper (Le *et al.* 2015) in the following manner: (1) it provides a more thorough description of the ASP-DPOP algorithm; (2) it elaborates on the algorithm’s theoretical properties with complete proofs; and (3) it includes additional experimental results.

† This research is partially supported by NSF grants HRD-1345232 and DGE-0947465.

the distributed management of power distribution networks (Kumar *et al.* 2009; Jain *et al.* 2012), the distributed generation of coalition structures (Ueda *et al.* 2010) and the distributed coordination of logistics operations (Léauté and Faltings 2011).

The field has matured considerably over the past decade, since the seminal ADOPT paper (Modi *et al.* 2005), as researchers continue to develop more sophisticated solving algorithms. The majority of the DCOP resolution algorithms can be classified in one of three classes: (1) *Search-based algorithms*, like ADOPT (Modi *et al.* 2005) and its variants (Yeoh *et al.* 2009, 2010; Gutierrez *et al.* 2011, 2013), Asynchronous Forward-Bounding (AFB) (Gershman *et al.* 2009), and MGM (Maheswaran *et al.* 2004), where the agents enumerate combinations of value assignments in a decentralized manner; (2) *Inference-based algorithms*, like DPOP (Petcu and Faltings 2005a) and its variants (Petcu and Faltings 2005b, 2007; Petcu *et al.* 2007, 2008), max-sum (Farinelli *et al.* 2008), and Action GDL (Vinyals *et al.* 2011), where the agents use dynamic programming techniques to propagate aggregated information to other agents; and (3) *Sampling-based algorithms*, like DUCT (Ottens *et al.* 2012) and D-Gibbs (Nguyen *et al.* 2013; Fioretto *et al.* 2014), where the agents sample the search space in a decentralized manner.

The existing algorithms have been designed and developed almost exclusively using *imperative programming* techniques, where the algorithms define a control flow, that is, a sequence of commands to be executed. In addition, the local solver employed by each agent is an “ad-hoc” implementation. In this paper, we are interested in investigating the benefits of using *declarative programming* techniques to solve DCOPs, along with the use of a general constraint solver, used as a black box, as each agent’s local constraint solver. Specifically, we propose an integration of *Distributed Pseudo-tree Optimization Procedure (DPOP)* (Petcu and Faltings 2005a), a popular DCOP algorithm, with *Answer Set Programming (ASP)* (Niemelä 1999; Marek and Truszczyński 1999) as the local constraint solver of each agent.

This paper provides the first step in bridging the areas of DCOPs and ASP; in the process, we offer novel contributions to both the DCOP field as well as the ASP field. *For the DCOP community*, we demonstrate that the use of ASP as a local constraint solver provides a number of benefits, including the ability to capitalize on (i) the highly expressive ASP language to more concisely define input instances (e.g., by representing constraint utilities as implicit functions instead of explicitly enumerating their extensions) and (ii) the highly optimized ASP solvers to exploit problem structure (e.g., propagating hard constraints to ensure consistency). *For the ASP community*, the paper makes the equally important contribution of increasing the applicability of ASP to model and solve a wide array of multi-agent coordination and resource allocation problems, currently modeled as DCOPs. Furthermore, it also demonstrates that general, off-the-shelf ASP solvers, which are continuously honed and improved, can be coupled with distributed message passing protocols to outperform specialized imperative solvers.

The paper is organized as follows. In Section 2, we review the basic definitions of DCOPs, the DPOP algorithm, and ASP. In Section 3, we describe in detail the structure of the novel ASP-based DCOP solver, called *ASP-DPOP*, and its implementation. Section 4 provides an analysis of the properties of ASP-DPOP,

including proofs of soundness and completeness of ASP-DPOP. Section 5 provides some experimental results, while Section 6 reviews related work. Finally, Section 7 provides conclusions and indications for future work.

## 2 Background

In this section, we present an overview of DCOPs, we describe DPOP, a complete distributed algorithm to solve DCOPs, and provide some fundamental definitions of ASP.

### 2.1 Distributed constraint optimization problems

A DCOP (Mailler and Lesser 2004; Modi *et al.* 2005; Petcu and Faltings 2005a; Yeoh and Yokoo 2012) can be described as a tuple  $\mathcal{M} = \langle \mathcal{X}, \mathcal{D}, \mathcal{F}, \mathcal{A}, \alpha \rangle$ , where

- $\mathcal{X} = \{x_1, \dots, x_n\}$  is a finite set of (decision) variables;
- $\mathcal{D} = \{D_1, \dots, D_n\}$  is a set of finite domains, where  $D_i$  is the domain of the variable  $x_i \in \mathcal{X}$ , for  $1 \leq i \leq n$ ;
- $\mathcal{F} = \{f_1, \dots, f_m\}$  is a finite set of constraints, where  $f_j$  is a  $k_j$ -ary function  $f_j : D_{j_1} \times D_{j_2} \times \dots \times D_{j_{k_j}} \mapsto \mathbb{R} \cup \{-\infty\}$  that specifies the utility of each combination of values of variables in its scope; the scope is denoted by  $scp(f_j) = \{x_{j_1}, \dots, x_{j_{k_j}}\}$ <sup>1</sup>;
- $\mathcal{A} = \{a_1, \dots, a_p\}$  is a finite set of agents; and
- $\alpha : \mathcal{X} \mapsto \mathcal{A}$  maps each variable to an agent.

We say that a variable  $x$  is owned by an agent  $a$  if  $\alpha(x) = a$ . We denote with  $\alpha_i$  the set of all variables that are owned by an agent  $a_i$ , i.e.,  $\alpha_i = \{x \in \mathcal{X} \mid \alpha(x) = a_i\}$ . Each constraint in  $\mathcal{F}$  can be either *hard*, indicating that some value combinations result in a utility of  $-\infty$  and must be avoided, or *soft*, indicating that all value combinations result in a finite utility and need not be avoided. A *value assignment* is a (partial or complete) function  $x$  that maps variables of  $\mathcal{X}$  to values in  $\mathcal{D}$ , such that if  $x(x_i)$  is defined, then  $x(x_i) \in D_i$  for  $i = 1, \dots, n$ . For the sake of simplicity, and with a slight abuse of notation, we will often denote  $x(x_i)$  simply with  $x_i$ . Given a constraint  $f_j$  and a *complete value assignment*  $x$  for all decision variables, we denote with  $x_{f_j}$  the projection of  $x$  to the variables in  $scp(f_j)$ ; we refer to this as a *partial value assignment* for  $f_j$ . For a DCOP  $\mathcal{M}$ , we denote with  $\mathcal{C}(\mathcal{M})$  the set of all complete value assignments for  $\mathcal{M}$ .

A *solution* of a DCOP is a complete value assignment  $\mathbf{x}$  for all variables such that

$$\mathbf{x} = \operatorname{argmax}_{\mathbf{x} \in \mathcal{C}(\mathcal{M})} \sum_{j=1}^m f_j(x_{f_j}) \quad (2.1)$$

A DCOP can be described by its *constraint graph*—i.e., a graph whose nodes correspond to agents in  $\mathcal{A}$  and whose edges connect pairs of agents who own variables in the scope of the same constraint.

<sup>1</sup> For the sake of simplicity, we assume a given ordering of variables.

*Definition 1 (Constraint Graph)*

A *constraint graph* of a DCOP  $\mathcal{M} = \langle \mathcal{X}, \mathcal{D}, \mathcal{F}, \mathcal{A}, \alpha \rangle$  is an undirected graph  $G_{\mathcal{M}} = (V, E)$ , where  $V = \mathcal{A}$  and

$$E = \{ \{a, a'\} \mid \{a, a'\} \subseteq \mathcal{A}, \exists f \in \mathcal{F} \text{ and } \{x_i, x_j\} \subseteq \mathcal{X}, \text{ such that } \{x_i, x_j\} \subseteq \text{scp}(f), \text{ and } \alpha(x_i) = a, \alpha(x_j) = a' \}. \tag{2.2}$$

Given the constraint graph  $G_{\mathcal{M}}$  and given a node  $a \in \mathcal{A}$ , we denote with  $N(a)$  the *neighbors* of  $a$ , i.e.,

$$N(a) = \{ a' \in \mathcal{A} \mid \{a, a'\} \in E \}. \tag{2.3}$$

*Definition 2 (Pseudo-tree)*

A *pseudo-tree* of a DCOP is a subgraph of  $G_{\mathcal{M}}$  that has the same nodes as  $G_{\mathcal{M}}$  such that (i) the included edges (called *tree edges*) form a rooted tree, and (ii) two nodes that are connected to each other in  $G_{\mathcal{M}}$  appear in the same branch of the tree.

The edges of  $G_{\mathcal{M}}$  that are not included in a pseudo-tree are called *back edges*. Notice that tree edges connect a node with its parent and its children, while back edges connect a node with its pseudo-parents and pseudo-children—i.e., nodes closer to the root are parents or pseudo-parents, while those closer to the leaves are children or pseudo-children. A pseudo-tree of a DCOP can be constructed using distributed DFS algorithms (Hamadi *et al.* 1998) applied to the constraint graph of the DCOP.

In this paper, we say that two variables are constrained to each other if they are in the scope of the same constraint. Given a pseudo-tree, the *separator* of a node  $a_i$  is, intuitively, the set of variables that (i) are owned by the ancestors of  $a_i$ , and (ii) are constrained with some variables that are either owned by  $a_i$  or the descendants of  $a_i$ . Formally, in a pseudo-tree, the *separator* of a node  $a_i$ , denoted by  $sep_i$ , is

$$sep_i = \{ x_{i'} \in \mathcal{X} \mid \alpha(x_{i'}) = a_{i'} \text{ where } a_{i'} \text{ is an ancestor of } a_i; \text{ and } \exists x_{i''} \in \mathcal{X}, f \in \mathcal{F}, \text{ such that } a_{i''} \text{ is either } a_i \text{ or a descendant of } a_i, \alpha(x_{i''}) = a_{i''}, \text{ and } \{x_{i'}, x_{i''}\} \subseteq \text{scp}(f) \} \tag{2.4}$$

We denote with  $P_i$ ,  $PP_i$ ,  $C_i$ , and  $PC_i$  the parent, the set of pseudo-parents, the set of children, and the set of pseudo-children of a node  $a_i$ , respectively. For simplicity, if  $A$  is a set of agents in  $\mathcal{A}$ , we also denote with  $\alpha_A$  the set of variables in  $\mathcal{X}$  that are owned by agents in  $A$ .

*Example 1*

Figure 1(a) shows the constraint graph of a DCOP  $\mathcal{M} = \langle \mathcal{X}, \mathcal{D}, \mathcal{F}, \mathcal{A}, \alpha \rangle$ , where

- $\mathcal{X} = \{x_1, x_2, x_3\}$ ;
- $\mathcal{D} = \{D_1, D_2, D_3\}$  where  $D_i = \{0, 1\}$  ( $1 \leq i \leq 3$ ) is the domain of the variable  $x_i \in \mathcal{X}$ ;
- $\mathcal{F} = \{x1\_cons\_x2, x1\_cons\_x3, x2\_cons\_x3\}$  where, for each  $1 \leq i < j \leq 3$ ,
  - for the constraint  $xi\_cons\_xj$  we have that  $\text{scp}(xi\_cons\_xj) = \{x_i, x_j\}$ ;
  - the utilities specified by the constraint  $xi\_cons\_xj$  are given in Figure 1(c).

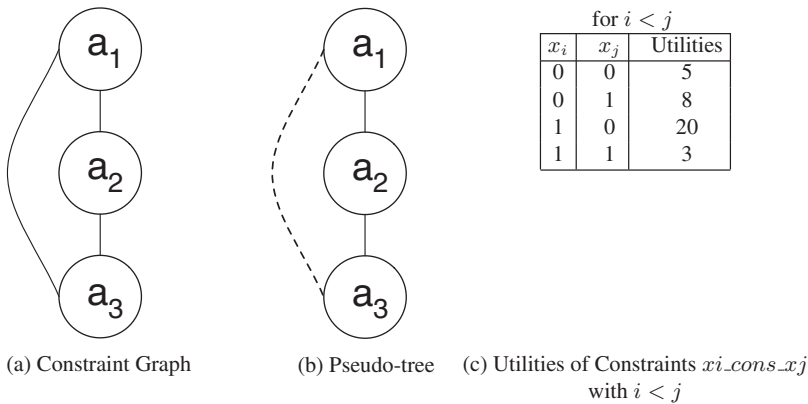


Fig. 1. Example DCOP. (a) Constraint graph. (b) Pseudo-tree. (c) Utilities of constraints  $x_{i\_cons\_xj}$  with  $i < j$ .

- $\mathcal{A} = \{a_1, a_2, a_3\}$ ; and
- $\alpha$  maps each variable  $x_i$  to agent  $a_i$ .

Figure 1(b) shows one possible pseudo-tree, where the dotted line is a back edge. In this pseudo-tree,  $P_3 = a_2$ ,  $PP_3 = \{a_1\}$ ,  $C_1 = \{a_2\}$ ,  $PC_1 = \{a_3\}$ , and  $sep_3 = \{x_1, x_2\}$ .

In a pseudo-tree  $T$  of a DCOP  $\langle \mathcal{X}, \mathcal{D}, \mathcal{F}, \mathcal{A}, \alpha \rangle$ , given  $a_i \in \mathcal{A}$  let  $R_{a_i}^T$  be the set of constraints in  $\mathcal{F}$ , such that

$$R_{a_i}^T = \{f \in \mathcal{F} \mid scp(f) \subseteq \alpha_i \cup \alpha_{P_i} \cup \alpha_{PP_i} \wedge scp(f) \cap \alpha_i \neq \emptyset\} \quad (2.5)$$

In the following, without causing any confusion, we often omit the superscript in  $R_{a_i}^T$  (i.e.,  $R_{a_i}$ ) if there is only one pseudo-tree mentioned in the context.

### Example 2

Considering again the DCOP in Example 1 and its pseudo-tree in Figure 1(b), we have  $R_{a_3} = \{x_{1\_cons\_x3}, x_{2\_cons\_x3}\}$ .

## 2.2 The distributed Pseudo-tree optimization procedure

The *DPOP* (Petcu and Faltings 2005a) is a complete algorithm to solve DCOPs with the following three phases:<sup>2</sup> Pseudo-tree generation, UTIL propagation, and VALUE propagation.

### 2.2.1 Phase 1: Pseudo-tree generation phase

DPOP does not require the use of any specific algorithm to construct the pseudo-tree. However, in many implementations of DPOP, including those within the DCOPolis (Sultanik et al. 2007) and FRODO (Léauté et al. 2009) repositories,

<sup>2</sup> Here, we detail an extended version of DPOP described in Petcu and Faltings (2005a) that removes the assumption that each agent owns exactly one variable.

greedy approaches such as the Distributed DFS algorithm (Hamadi *et al.* 1998) are used.

The distributed DFS algorithm operates as follows. First of all, the algorithm assigns a score to each agent, according to a heuristic function. It then selects an agent with the largest score as the root of the pseudo-tree. Once the root is selected, the algorithm initiates a DFS-traversal of the constraint graph, greedily adding the neighboring agent with the largest score as the child of the current agent. This process is repeated until all agents in the constraint graph are added to the pseudo-tree.

The agents' scores can be chosen arbitrarily. A commonly used heuristic is the max-degree heuristic  $h(a_i)$ :

$$h(a_i) = |N(a_i)| \quad (2.6)$$

which sets an agent's score to its number of neighbors. In situations where multiple agents have the same maximal score, the algorithm breaks ties according to a different heuristic, such as the variable-ID heuristic, which assigns to each agent a score that is equal to its unique ID. In our experiments, we use the max-degree heuristic and break ties with the variable-ID heuristic in the construction of the pseudo-tree.

### 2.2.2 Phase 2: UTIL propagation phase

The UTIL propagation phase is a bottom-up process, which starts from the leaves of the pseudo-tree and propagates upward, following only the tree edges of the pseudo-tree. In this process, the agents send UTIL messages to their parents.

*Definition 3 (UTIL Messages (Petcu 2009))*

$UTIL_{a_i}^{a_j}$ , the UTIL message sent by agent  $a_i$  to agent  $a_j$ , is a multi-dimensional matrix, with one dimension for each variable in  $sep_i$ . With a slight abuse of notation, we denote with  $scp(UTIL_{a_i}^{a_j})$  the set of variables in the message.

Instead of using a multi-dimensional matrix, one can also flatten the multi-dimensional matrix into a table where each row of the table is for one combination of value assignment of variables in  $sep_i$  and the respective utility for that combination. For simplicity, in this paper, we will represent UTIL messages under their tabular form. We can observe that it is always true that  $\alpha_j \cap scp(UTIL_{a_i}^{a_j}) \neq \emptyset$ . The semantics of such a UTIL message is similar to a constraint whose scope is the set of all variables in the context of the message (its dimensions). The size of such a UTIL message is the product of the domain sizes of variables in the context of the message.

Intuitively, a UTIL message summarizes the optimal sum of utilities in its subtree for each value combination of variables in its separator. An agent  $a_i$  computes its UTIL message by (i) summing the utilities in the UTIL messages received from its child agents and the utilities of constraints whose scopes are exclusively composed of the variables of  $a_i$  and the variables in its separator (i.e.,  $R_{a_i}$ ), and then (ii) projecting out the variables of  $a_i$ , by optimizing over them. Algorithm 1 provides a formal description of Phase 2.

**Algorithm 1: DPOP Phase 2 (UTIL Propagation Phase)**

```

1 Each agent  $a_i$  does:
2  $JOIN_{a_i}^{P_i} = null$ 
3 forall the  $a_c \in C_i$  do
4   |   wait for  $UTIL_{a_c}^{a_i}$  message to arrive from  $a_c$ 
5   |    $JOIN_{a_i}^{P_i} = JOIN_{a_i}^{P_i} \oplus UTIL_{a_c}^{a_i}$  // join UTIL messages from children as they arrive
6 end
7  $JOIN_{a_i}^{P_i} = JOIN_{a_i}^{P_i} \oplus (\oplus_{f \in R_{a_i}} f)$  // also join all constraints with parent/pseudo-parents
8  $UTIL_{a_i}^{P_i} = JOIN_{a_i}^{P_i} \perp_{\alpha_i}$  // use projection to eliminate its owned variables
9 Send  $UTIL_{a_i}^{P_i}$  message to its parent agent  $P_i$ 

```

Algorithm 1 uses the *JOIN* operator (i.e.,  $\oplus$ ) and the *PROJECTION* operator (i.e.,  $\perp$ ).

**Definition 4 (JOIN  $\oplus$  Operator)**

$U = UTIL_{a_k}^{a_i} \oplus UTIL_{a_l}^{a_i}$  is the join of two UTIL matrices (constraints).  $U$  is also a matrix (constraint) with  $scp(U) = scp(UTIL_{a_k}^{a_i}) \cup scp(UTIL_{a_l}^{a_i})$  as dimensions. For each possible combination  $x$  of values of variables in  $scp(U)$ , the corresponding value of  $U(x)$  is the sum of the corresponding cells in the two source matrices, i.e.,  $U(x) = UTIL_{a_k}^{a_i}(x_{UTIL_{a_k}^{a_i}}) + UTIL_{a_l}^{a_i}(x_{UTIL_{a_l}^{a_i}})$  where  $x_{UTIL_{a_k}^{a_i}}$  and  $x_{UTIL_{a_l}^{a_i}}$  are partial value assignments from  $x$  for all variables in  $scp(UTIL_{a_k}^{a_i})$  and  $scp(UTIL_{a_l}^{a_i})$ , respectively.

Since UTIL messages can be seen as constraints, the  $\oplus$  operator can be used to join UTIL messages and constraints.

**Example 3**

Given 2 constraints  $x1\_cons\_x3$  and  $x2\_cons\_x3$  in Example 1, let  $JOIN_{a_3}^{a_2} = x1\_cons\_x3 \oplus x2\_cons\_x3$ . It is possible to see that  $scp(JOIN_{a_3}^{a_2}) = \{x_1, x_2, x_3\}$ . The utility corresponding to  $x_1 = x_2 = x_3 = 0$  is  $JOIN_{a_3}^{a_2}(x_1 = 0, x_2 = 0, x_3 = 0) = 5 + 5 = 10$ . Moreover,  $JOIN_{a_3}^{a_2}(x_1 = 0, x_2 = 0, x_3 = 1) = 8 + 8 = 16$ .

For the  $\perp$  operator, knowing that  $\alpha_i \subseteq scp(JOIN_{a_i}^{P_i})$ ,  $JOIN_{a_i}^{P_i} \perp_{\alpha_i}$  is the projection through optimization of the  $JOIN_{a_i}^{P_i}$  matrix along axes representing variables in  $\alpha_i$ .

**Definition 5 (PROJECTION  $\perp$  Operator)**

Let  $\alpha_i$  be a set of variables where  $\alpha_i \subseteq scp(JOIN_{a_i}^{P_i})$ , and let  $X_i$  be the set of all possible value combinations of variables in  $\alpha_i$ . A matrix  $U = JOIN_{a_i}^{P_i} \perp_{\alpha_i}$  is defined as: (i)  $scp(U) = scp(JOIN_{a_i}^{P_i}) \setminus \alpha_i$ , and (ii) for each possible value combination  $x$  of variables in  $scp(U)$ ,  $U(x) = \max_{x' \in X_i} JOIN_{a_i}^{P_i}(x, x')$ .

**Example 4**

Considering again  $JOIN_{a_3}^{a_2}$  in Example 3, let  $U = JOIN_{a_3}^{a_2} \perp_{\{x_3\}}$ . We have  $scp(U) = \{x_1, x_2\}$ , and  $U(x_1 = 0, x_2 = 0) = \max(JOIN_{a_3}^{a_2}(x_1 = 0, x_2 = 0, x_3 = 0), JOIN_{a_3}^{a_2}(x_1 = 0, x_2 = 0, x_3 = 1)) = \max(10, 16) = 16$ .

As an example for the computations in Phase 2 (UTIL propagation phase), we consider again the DCOP in Example 1.

Table 1. UTIL phase computations in DPOP

| $x_1$ | $x_2$ | Utilities                   |
|-------|-------|-----------------------------|
| 0     | 0     | $\max(5 + 5, 8 + 8) = 16$   |
| 0     | 1     | $\max(5 + 20, 8 + 3) = 25$  |
| 1     | 0     | $\max(20 + 5, 3 + 8) = 25$  |
| 1     | 1     | $\max(20 + 20, 3 + 3) = 40$ |

(a)

| $x_1$ | Utilities                    |
|-------|------------------------------|
| 0     | $\max(5 + 16, 8 + 25) = 33$  |
| 1     | $\max(20 + 25, 3 + 40) = 45$ |

(b)

*Example 5*

In the DCOP in Example 1, the agent  $a_3$  computes its UTIL message,  $UTIL_{a_3}^{a_2}$  (see Table 1(a)), and sends it to its parent agent  $a_2$ . The agent  $a_2$  then computes its UTIL message,  $UTIL_{a_2}^{a_1}$  (see Table 1(b)), and sends it to its parent agent  $a_1$ . Finally, the agent  $a_1$  computes the optimal utility of the entire problem, which is 45.

## 2.2.3 Phase 3: VALUE propagation phase

Phase 2 finishes when the UTIL message reaches the root of the tree. At that point, each agent, starting from the root of the pseudo-tree, determines the optimal value for its variables based on (i) the computation from Phase 2, and (ii) (for non-root agent only) the VALUE message that is received from its parent. Then, it sends these optimal values to its child agents through VALUE messages. Algorithm 2 provides a formal description of Phase 3.

A VALUE message that travels from the parent  $P_i$  to the agent  $a_i$ ,  $VALUE_{P_i}^{a_i}$ , contains the optimal value assignment for variables owned by either the parent agent or the pseudo-parent agents of the agent  $a_i$ .

**Algorithm 2:** DPOP Phase 3 (VALUE Propagation Phase)

```

1 Each agent  $a_i$  do:
2 wait for  $VALUE_{P_i}^{a_i}(sep_i^*)$  message from its parent agent  $P_i$  //  $sep_i^*$  is the optimal value
  assignment for all variables in  $sep_i$ 
3  $\alpha_i^* \leftarrow \operatorname{argmax}_{\alpha_i \in X_i} JOIN_{a_i}^{P_i}(sep_i^*, \alpha_i)$  //  $X_i$  is the set of all possible value combinations of
  variables in  $\alpha_i$ 
4 forall the  $a_c \in C_i$  do
5   | let  $sep_i^{**}$  be the partial optimal value assignment for variables in  $sep_c$  from  $sep_i^*$ 
6   | send  $VALUE(sep_i^{**}, \alpha_i^*)$  as  $VALUE_{a_i}^{a_c}$  message to its child agent  $a_c$ 
7 end

```

*Example 6*

In the DCOP in Example 1, the agent  $a_1$  determines that the value with the largest utility for its variable  $x_1$  is 1, with a utility of 45, and then sends this information down to its child agent  $a_2$  in a VALUE message, i.e.,  $VALUE_{a_1}^{a_2}(x_1 = 1)$ . Upon receiving that VALUE message, the agent  $a_2$  determines that the value for its variable  $x_2$  with the largest utility of the subtree rooted at the agent  $a_2$ , assuming that  $x_1 = 1$ , is 0, with a utility of 45. The agent  $a_2$  then sends this information down



to its child agent  $a_3$ , i.e.,  $VALUE_{a_2}^{a_3}(x_1 = 1, x_2 = 0)$ . Finally, upon receiving such VALUE message, the agent  $a_3$  determines that the value for its variable  $x_3$  with the largest utility of the subtree rooted at the agent  $a_3$ , assuming that  $x_1 = 1$  and  $x_2 = 0$ , is 0, with a utility of 25.

### 2.3 Answer set programming

Let us provide some general background on ASP (see, for example, Baral 2003; Gelfond and Kahl 2006 for more details).

An *answer set program*  $\Pi$  is a set of *rules* of the form

$$c \leftarrow a_1, \dots, a_j, \text{not } a_{j+1}, \dots, \text{not } a_m \quad (2.7)$$

where  $0 \leq j \leq m$ , for  $1 \leq i \leq m$  each  $a_i$  or  $c$  is a literal of a first-order language  $\mathcal{L}$ , and *not* represents *negation-as-failure* (*naf*). For a literal  $a$ , *not a* is called a *naf-literal*. For a rule of the form (2.7), the left- and right-hand sides of the rule are called the *head* and the *body* of the rule, respectively. Both the head and the body can be empty. When the head is empty, the rule is called a *constraint*. When the body is empty, the rule is called a *fact*. A literal (resp. rule) is a *ground* literal (resp. ground rule) if it does not contain any variable. A rule with variables is simply used as a shorthand for the set of its ground instances from the language  $\mathcal{L}$ . Similarly, a non-ground program (i.e., a program containing some non-ground rules) is a shorthand for all ground instances of its rules. Throughout this paper, we follow the traditional notation in writing ASP rules, where names that start with an upper case letter represent variables. For a ground instance  $r$  of a rule of the form (2.7),  $head(r)$  denotes the set  $\{c\}$ , while  $pos(r)$  and  $neg(r)$  denote  $\{a_1, \dots, a_j\}$  and  $\{a_{j+1}, \dots, a_m\}$ , respectively.

Let  $X$  be a set of ground literals.  $X$  is consistent if there is no atom  $a$  such that  $\{a, \neg a\} \subseteq X$ . The body of a ground rule  $r$  of the form (2.7) is *satisfied* by  $X$  if  $neg(r) \cap X = \emptyset$  and  $pos(r) \subseteq X$ . A ground rule of the form (2.7) with nonempty head is satisfied by  $X$  if either its body is not satisfied by  $X$  or  $head(r) \cap X \neq \emptyset$ . A constraint is *satisfied* by  $X$  if its body is not satisfied by  $X$ .

For a consistent set of ground literals  $S$  and a ground program  $\Pi$ , the *reduct* of  $\Pi$  w.r.t.  $S$ , denoted by  $\Pi^S$ , is the program obtained from  $\Pi$  by deleting (i) each rule that has a naf-literal *not a* in its body where  $a \in S$ , and (ii) all naf-literals in the bodies of the remaining rules.

$S$  is an *answer set* (or a *stable model*) of a ground program  $\Pi$  (Gelfond and Lifschitz 1990) if it satisfies the following conditions: (i) if  $\Pi$  does not contain any naf-literal (i.e.,  $j = m$  in every rule of  $\Pi$ ), then  $S$  is a minimal consistent set of literals that satisfies all the rules in  $\Pi$ ; and (ii) if  $\Pi$  contains some naf-literals ( $j < m$  in some rules of  $\Pi$ ), then  $S$  is an answer set of  $\Pi^S$ . Note that  $\Pi^S$  does not contain naf-literals, and thus its answer set is defined in case (i). A program  $\Pi$  is said to be *consistent* if it has some answer sets. Otherwise, it is inconsistent.

The ASP language includes also language-level extensions to facilitate the encoding of aggregates (*min*, *max*, *sum*, etc.).

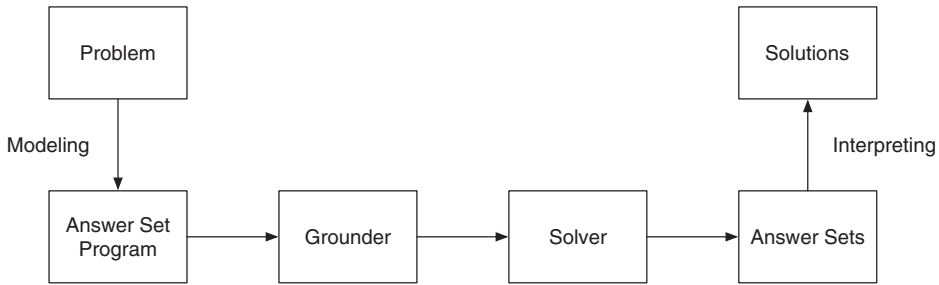


Fig. 2. Solving a problem using ASP.

*Example 7*

Let us consider an ASP program  $\Pi$  that consists of two facts and one rule:

$$\text{int}(5) \leftarrow \quad (2.8)$$

$$\text{int}(10) \leftarrow \quad (2.9)$$

$$\text{max}(U) \leftarrow U = \# \text{max}\{V : \text{int}(V)\} \quad (2.10)$$

The third rule uses an aggregate to determine the maximum in the set  $\{V : \text{int}(V)\}$ .  $\Pi$  has one answer set:  $\{\text{int}(5), \text{int}(10), \text{max}(10)\}$ . Thus,  $\Pi$  is consistent.

Moreover, to increase the expressiveness of logic programming and simplify its use in applications, the syntax of ASP has been extended with *choice rules*. Choice rules are of the form:

$$l\{a_1, \dots, a_m\}u \leftarrow a_{m+1}, \dots, a_n, \text{not } a_{n+1}, \dots, \text{not } a_k \quad (2.11)$$

where  $l\{a_1, \dots, a_m\}u$  is called a choice atom,  $l$  and  $u$  are integers,  $l \leq u$ ,  $0 \leq m \leq n \leq k$ , and each  $a_i$  is a literal for  $1 \leq i \leq k$ . This rule allows us to derive any subset of  $\{a_1, \dots, a_m\}$  whose cardinality is between the lower bound  $l$  and upper bound  $u$  whenever the body is satisfied.  $l$  or  $u$  can be omitted. If  $l$  is omitted,  $l = 0$ , and if  $u$  is omitted,  $u = +\infty$ . Standard syntax for choice rules has been proposed and adopted in most state-of-the-art ASP solvers, such as CLASP (Gebser *et al.* 2007) and DLV (Citrigno *et al.* 1997).

Figure 2 visualizes how to solve a problem using ASP. In more detail, the problem is encoded as an answer set program whose answer sets correspond to solutions. The answer set program, which may contain variables, is then grounded using an ASP grounder, e.g., GRINGO (Gebser *et al.* 2011). The grounding process employs smart techniques to reduce the size of the resulting ground program, e.g., removing literals from rules that are known to be true, removing rules that will not contribute to the computation of answer sets.

*Example 8*

Let us consider an ASP program  $\Pi$  that consists of two facts and one rule:

$$\text{int}(1) \leftarrow \quad (2.12)$$

$$\text{int}(-1) \leftarrow \quad (2.13)$$

$$\text{isPositive}(X) \leftarrow \text{int}(X), X > 0 \quad (2.14)$$

Using a naive grounder that simply replaces consistently the variable  $X$  with the two constants 1 and  $-1$ , the ground program of  $\Pi$  consists of the two facts (2.12) and (2.13) and the two following ground rules:

$$\text{isPositive}(1) \leftarrow \text{int}(1), 1 > 0 \quad (2.15)$$

$$\text{isPositive}(-1) \leftarrow \text{int}(-1), -1 > 0 \quad (2.16)$$

It is easy to see that the ground rule (2.16) is unnecessary (i.e., its body cannot be satisfied by any set of literals due to the literal  $-1 > 0$ ) and should be removed. In contrast, the ground program of  $\Pi$  obtained by GRINGO has only three facts: (2.12), (2.13), and

$$\text{isPositive}(1) \leftarrow \quad (2.17)$$

We observe that the unnecessary rule (2.16) is omitted since its body cannot be satisfied (i.e.,  $-1 > 0$ ), and the fact (2.17) is obtained from the rule (2.15) by removing all literals in its body because the grounder can determine as been always satisfied.

All the answer sets of the program produced by the ASP grounder are then computed by an ASP solver, e.g., CLASP (Gebser et al. 2007). The solutions to the original problem can be determined by properly interpreting the different answer sets computed, where each answer sets corresponds to one of the possible solutions to the original problem. For readers who are interested in how to solve an answer set program, the foundations and algorithms underlying the grounding and solving technology used in GRINGO and CLASP are described in detail in Gebser et al. (2012) and Kaufmann et al. (2016).

### 3 ASP-DPOP

*ASP-DPOP* is a framework that uses logic programming to capture the structure of DCOPs, and to emulate the computation and communication operations of DPOP. In particular, each agent in a DCOP is represented by a separate ASP program—effectively enabling the infusion of a knowledge representation framework in the DCOP paradigm.

The overall communication infrastructure required by the distributed computation of DPOP is expressed using a subset of the *SICStus Prolog* language (Carlsson et al. 2015), extended with multi-threading and the *Linda* blackboard facilities. In ASP-DPOP, we use CLASP (Gebser et al. 2007), with its companion grounder GRINGO, as our ASP solver, being the current state-of-the-art for ASP. In this section, we will describe the structure of ASP-DPOP and its implementation.

#### 3.1 The architecture of ASP-DPOP

ASP-DPOP is an agent architecture that reflects the structure of DCOPs, where several agents reflect the computation and communication operations of DPOP. The internal structure of each *ASP-DPOP agent*, shown in Figure 3, is composed of




Fig. 3. The structure of an ASP-DPOP agent.

two modules. The first module is the *Specification Module (SM)*, which encloses an ASP program that captures a corresponding agent as specified in the DCOP—i.e., the agent’s name, the agent’s neighbors, the description of the variables owned by the agent, the description of the variables owned by the agent’s neighbors, and the description of the constraints whose scope include any of the variables owned by the agent.

The second module is the *Controller Module (CM)*, encoded as a Prolog program. The CM instructs the agent to perform the communication operations of DPOP, such as cooperating with other agents to generate a pseudo-tree, waiting for UTIL messages from child agents, sending the UTIL message to the parent agent (if present), waiting for the VALUE message from the parent agent (if present), and sending the VALUE messages to the child agents.

In ASP-DPOP, each DCOP is represented by a set of ASP-DPOP agents; each agent is modeled by its knowledge bases, located at its SM and CM, and it interacts with other agents in accordance to the rules of its CM.

### 3.2 ASP-DPOP implementation: Specification module (SM)

Let us describe how to capture the structure of a DCOP in the SM of an ASP-DPOP agent using ASP. Let us consider a generic DCOP  $\mathcal{M} = \langle \mathcal{X}, \mathcal{D}, \mathcal{F}, \mathcal{A}, \alpha \rangle$ . We represent  $\mathcal{M}$  using a set of ASP-DPOP agents whose SMs are ASP programs  $\{\Pi_{a_i} \mid a_i \in \mathcal{A}\}$ . We will show how to generate  $\Pi_{a_i}$  for each agent  $a_i$ . In the following, we say  $a$  and  $a'$  in  $\mathcal{A}$  are *neighbors* if there exists  $x$  and  $x'$  in  $\mathcal{X}$  such that  $\alpha(x) = a$ ,  $\alpha(x') = a'$ , and there is a  $f \in \mathcal{F}$  such that  $\{x, x'\} \subseteq \text{scp}(f)$ . Given a constraint  $f \in \mathcal{F}$ , we say that  $f$  is *owned by the agent*  $a_i$  if the scope of  $f$  contains some variables owned by the agent  $a_i$ .<sup>3</sup>

For each variable  $x_i \in \mathcal{X}$  we define a collection  $L(x_i)$  of ASP rules that includes the following:

<sup>3</sup> The concept of *ownership of a constraint* is introduced to facilitate the representation of ASP-DPOP implementation. Intuitively, an agent should know about a constraint if the agent owns some variables that are in the scope of such constraint. Under this perspective, a constraint may be owned by several agents.

- A fact of the form

$$\text{variable}(x_i) \leftarrow \quad (3.1)$$

for identifying the name of the variable;

- For each  $d \in D_i \in \mathcal{D}$ , a fact of the form

$$\text{value}(x_i, d) \leftarrow \quad (3.2)$$

for identifying the possible values of  $x_i$ . Alternatively, if the domain  $D_i$  is an integer interval  $[\text{lower\_bound} \dots \text{upper\_bound}]$ , we can use the additional facts of the form

$$\text{begin}(x_i, \text{lower\_bound}) \leftarrow \quad (3.3)$$

$$\text{end}(x_i, \text{upper\_bound}) \leftarrow \quad (3.4)$$

to facilitate the description of the domain  $D_i$ . In such case, the *value* predicates similar to ones in (3.2) are achieved by the rule

$$\text{value}(X, B..E) \leftarrow \text{variable}(X), \text{begin}(X, B), \text{end}(X, E) \quad (3.5)$$

Intuitively,  $B$  and  $E$  in (3.5) are variables that should be grounded with *lower\_bound* and *upper\_bound* from (3.3)–(3.4), respectively.

For each constraint  $f_j \in \mathcal{F}$ , where  $\text{scp}(f_j) = \{x_{j_1}, \dots, x_{j_{k_j}}\}$ , we define a collection  $L(f_j)$  of rules that includes the following:

- A fact of the form

$$\text{constraint}(f_j) \leftarrow \quad (3.6)$$

for identifying the name of the constraint;

- For each variable  $x \in \text{scp}(f_j)$ , a fact of the form

$$\text{scope}(f_j, x) \leftarrow \quad (3.7)$$

for identifying the scope of the constraint; and

- For each partial value assignment  $x_{f_j}$  for all variables in  $\text{scp}(f_j)$ , where  $v_{j_1}, \dots, v_{j_{k_j}}$  are the value assignments of the variables  $x_{j_1}, \dots, x_{j_{k_j}}$ , respectively, such that  $f_j(x_{f_j}) = u \neq -\infty$ , a fact of the form

$$f_j(u, v_{j_1}, \dots, v_{j_{k_j}}) \leftarrow \quad (3.8)$$

For each partial value assignment  $x_{f_j}$  for all variables in  $\text{scp}(f_j)$ , where  $v_{j_1}, \dots, v_{j_{k_j}}$  are the value assignments of the variables  $x_{j_1}, \dots, x_{j_{k_j}}$ , respectively, such that  $f_j(x_{f_j}) = -\infty$ , a fact of the form<sup>4</sup>

$$f_j(\#inf, v_{j_1}, \dots, v_{j_{k_j}}) \leftarrow \quad (3.9)$$

<sup>4</sup> *#inf* is a special constant representing the smallest possible value in ASP language.

Alternatively, it is also possible to envision cases where the utility of a constraint is implicitly modeled by logic programming rules, as shown in the following example. It is important to mention that, considering a constraint  $f_j \in \mathcal{F}$ :

- (1) the order of variables (e.g.,  $x_{j_1}, \dots, x_{j_{k_j}}$ ) in  $scp(f_j)$ , whose corresponding value assignments (e.g.,  $v_{j_1}, \dots, v_{j_{k_j}}$ ) that appear in facts of the forms (3.8) and (3.9), needs to be consistent in all facts of the forms (3.8) and (3.9) that relate to the constraint  $f_j$ ; and
- (2) the order of the facts of the form (3.7) that are added to  $L(f_j)$  to identify the scope of the constraint  $f_j$  needs to be consistent with the order of variables (e.g.,  $x_{j_1}, \dots, x_{j_{k_j}}$ ) mentioned in (1).

These requirements (i.e., (1) and (2)) are critical, because they allow CM to understand which variables belong to what values that appear in the facts of the forms (3.8) and (3.9), when CMs read  $L(f_j)$ . This is done because, in *SICStus Prolog*, the search rule is “search forward from the beginning of the program.” Therefore, the order of the predicates (i.e., facts) that are added to *SICStus Prolog* matters.

#### Example 9

Let us consider a constraint  $f$  whose scope is  $\{x, x'\}$ , and  $f$  specifies that the utility of value assignments  $x = v$ ,  $x' = v'$  is  $v + v'$ . The facts of the form (3.8) for the constraint  $f$  can be modeled by the following rule:

$$f(V + V', V, V') \leftarrow value(x, V), value(x', V') \quad (3.10)$$

For each agent  $a_i$  we define an ASP program  $\Pi_{a_i}$  that includes the following:

- A fact of the form

$$agent(a_i) \leftarrow \quad (3.11)$$

for identifying the name of the agent;

- For each variable  $x \in \mathcal{X}$  that is owned by the agent  $a_i$ , a fact of the form

$$owner(a_i, x) \leftarrow \quad (3.12)$$

- For each agent  $a_j$  who is a neighbor of the agent  $a_i$ , a fact of the form

$$neighbor(a_j) \leftarrow \quad (3.13)$$

- For each variable  $x' \in \mathcal{X}$  that is owned by an agent  $a_j$  who is a neighbor of the agent  $a_i$ , a fact of the form

$$owner(a_j, x') \leftarrow \quad (3.14)$$

- For each constraint  $f_j \in \mathcal{F}$  owned by the agent  $a_i$ , the set of rules

$$L(f_j) \quad (3.15)$$

- For each variable  $x \in \mathcal{X}$  that is in the scope of some constraints owned by the agent  $a_i$ , the set of rules

$$L(x) \tag{3.16}$$

### 3.3 ASP-DPOP implementation: Encoding UTIL and VALUE messages

The ASP-DPOP framework emulates the computation and communication operations of DPOP, where each ASP-DPOP agent produces UTIL and VALUE messages, and forwards them to its parent and child agents, as DPOP does. In ASP-DPOP, UTIL and VALUE messages are encoded as ASP facts, as discussed in this subsection.

#### 3.3.1 UTIL messages

In DPOP, each UTIL message sent from a child agent  $a_i$  to its parent agent  $P_i$  is a matrix. In encoding a UTIL message in ASP-DPOP, we represent each cell of the matrix of the UTIL message, whose associated utility is not  $-\infty$ , as an ASP atom of the form:

$$table\_max\_a_i(u, v_{i_1}, \dots, v_{i_{k_i}}) \tag{3.17}$$

which indicates that the optimal aggregate utility of the value assignments  $x_{i_1} = v_{i_1}, \dots, x_{i_{k_i}} = v_{i_{k_i}}$  is  $u \neq -\infty$ , where  $sep_i = \{x_{i_1}, \dots, x_{i_{k_i}}\}$ . In other words, the parent agent  $P_i$  knows that  $UTIL_{a_i}^{P_i}(x_{i_1} = v_{i_1}, \dots, x_{i_{k_i}} = v_{i_{k_i}}) = u \neq -\infty$  if it receives the fact  $table\_max\_a_i(u, v_{i_1}, \dots, v_{i_{k_i}})$ . It is important to know that the encoding of a UTIL message omits the cells whose associated utilities are  $-\infty$ .

In addition to facts of the form (3.17),  $a_i$  also informs  $P_i$  about variables in its separator. Thus, the encoding of the UTIL message sent from the agent  $a_i$  to the agent  $P_i$  includes also atoms of the form:

$$table\_info(a_i, a_{i_1}, x_{i_1}, lb_{i_1}, ub_{i_1}) \tag{3.18}$$

...

$$table\_info(a_i, a_{i_{k_i}}, x_{i_{k_i}}, lb_{i_{k_i}}, ub_{i_{k_i}}) \tag{3.19}$$

Each fact  $table\_info(a_i, a_{i_t}, x_{i_t}, lb_{i_t}, ub_{i_t})$  informs  $P_i$  that  $x_{i_t}$  is a variable in the separator of  $a_i$  whose domain is specified by  $lb_{i_t}$  (lower bound) and  $ub_{i_t}$  (upper bound) and whose owner is  $a_{i_t}$ .<sup>5</sup> It is also critical to note that the order of the atoms of the forms (3.18), (3.19) matters, since such order will allow the respective CM understand which variable belongs to the values stated in facts of the form (3.17) after reading such encoded UTIL messages.

<sup>5</sup> For simplicity, we assume that the domains  $D_i$  are integer intervals.

*Example 10*

Consider again the DCOP in Example 1. The UTIL message, sent from the agent  $a_3$  to the agent  $a_2$ , in Table 1(a) is encoded as the set of the ASP atoms:

$$\text{table\_max\_}a_3(16, 0, 0) \quad (3.20)$$

$$\text{table\_max\_}a_3(25, 0, 1) \quad (3.21)$$

$$\text{table\_max\_}a_3(25, 1, 0) \quad (3.22)$$

$$\text{table\_max\_}a_3(40, 1, 1) \quad (3.23)$$

$$\text{table\_info}(a_3, a_1, x_1, 0, 1) \quad (3.24)$$

$$\text{table\_info}(a_3, a_2, x_2, 0, 1) \quad (3.25)$$

*Example 11*

Similarly, considering again the DCOP in Example 1, the UTIL message sent from the agent  $a_2$  to the agent  $a_1$  in Table 1(b) is encoded as the set of ASP facts:

$$\text{table\_max\_}a_2(33, 0) \quad (3.26)$$

$$\text{table\_max\_}a_2(45, 1) \quad (3.27)$$

$$\text{table\_info}(a_2, a_1, x_1, 0, 1) \quad (3.28)$$

### 3.3.2 VALUE messages

In DPOP, each VALUE message sent from a parent agent  $P_i$  to its child agents  $a_i$  contains the optimal value assignment for variables owned by either the parent agent or the pseudo-parent agents of the agent  $a_i$ . Thus, in encoding a VALUE message, we use atoms of the form:

$$\text{solution}(a, x, v) \quad (3.29)$$

where  $v$  is the value assignment of the variable  $x$  owned by the agent  $a$  for an optimal solution.

*Example 12*

Consider again the DCOP in Example 1. The VALUE message sent from the agent  $a_1$  to the agent  $a_2$  is encoded as the ASP atom:

$$\text{solution}(a_1, x_1, 1) \quad (3.30)$$

Similarly, the VALUE message sent from the agent  $a_2$  to the agent  $a_3$  is encoded as the set of the ASP atoms:

$$\text{solution}(a_1, x_1, 1) \quad (3.31)$$

$$\text{solution}(a_2, x_2, 0) \quad (3.32)$$

### 3.4 ASP-DPOP implementation: Controller module (CM)

The CM in each ASP-DPOP agent  $a_i$ , denoted by  $C_{a_i}$ , consists of a set of Prolog rules for communication (sending, receiving, and interpreting messages) and a set of



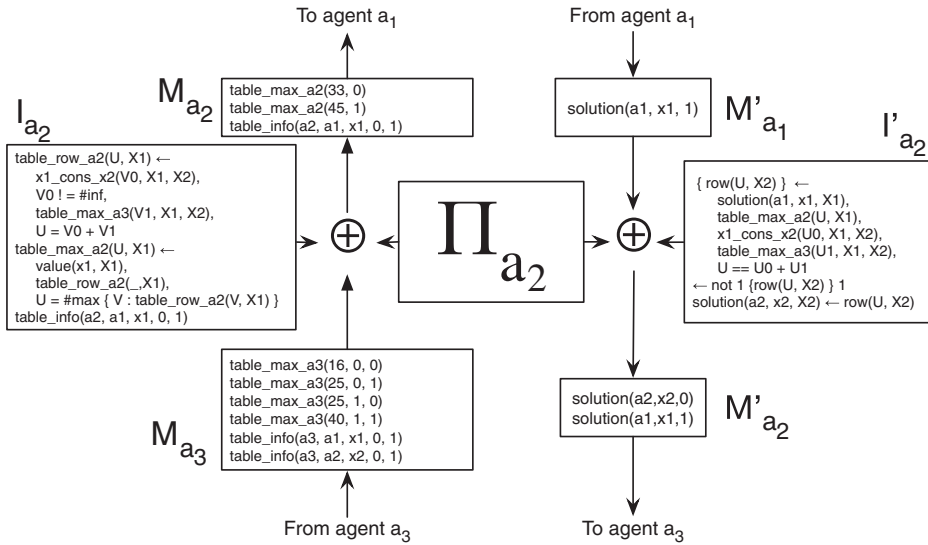


Fig. 4. Phase 2 and Phase 3 of agent  $a_2$  in ASP-DPOP on DCOP in Example 1.

rules for generating an ASP program that is used for the computations of a UTIL message and a VALUE message. In this subsection, we would like to discuss some code fragments to show how  $C_{a_i}$  is structured.<sup>6</sup> To begin with, we will show how  $C_{a_i}$  uses the Linda blackboard library of Prolog to exchange the messages.

There are three types of messages exchanged through the Linda blackboard; they are *tree*, *util*, and *value* messages that are used in Phase 1, Phase 2, and Phase 3, respectively, of DPOP. For sending (resp. waiting for) a message, we use the built-in predicate *out/1* (resp. *in/1*) provided by the Linda library of Prolog. Every message is formatted as *message(From, To, Type, Data)* where the arguments denote the agent who sends this message, the agent who should receive this message, the type of the message, and the data enclosed in the message, respectively. The implementation of the communication and the three phases of DPOP are described next.

### 3.4.1 Sending messages

The following Prolog rule generates a message of type  $t \in \{tree, util, value\}$ , with content  $d$  (Content), to be sent from an agent  $a_i$  (From) to an agent  $a_k$  (To):

```
% sending message
send_message(From, To, Type, Content) :-
    out(message(From, To, Type, Content)).
```

### 3.4.2 Waiting for messages

The following Prolog rule instructs agent  $a_k$  ( $a\_k$ ) to wait for a message:

<sup>6</sup> The code listed in this section is a simplified version of the actual code for  $C_{a_i}$ , showing a condensed version of the clause bodies; however, it still gives a flavor of the implementation of  $C_{a_i}$  and should be sufficiently explanatory for the purpose of the CM.

```
% waiting for a message
wait_message(From, a_k, Type, Data) :- in(message(From, a_k, Type, Data)).
```

where *From*, *Type*, and *Data* will be unified with the name of the agent who sent this message, the type of the message, and the data enclosed in the message, respectively.

### 3.4.3 Creating the pseudo-tree: Phase 1

In this phase, ASP-DPOP agents cooperate with each other to construct a pseudo-tree. For simplicity, we will show here the clauses in  $C_{a_i}$  for generating a pseudo-tree by initiating a DFS-traversal. We assume that the agent  $a_i$  is not the root of the pseudo-tree. The agent  $a_i$  waits for a *tree* message from an agent *Parent*. The content (*Data*) enclosed in such a *tree* message is the set of *visited* agents—i.e., the agents who have already started performing the DFS. Upon receiving a *tree* message,  $a_i$  will execute the following clauses:

```
% pseudo-tree generation
generate_tree(Parent, Data) :-
    assign_parent(Parent),
    assign_pseudo_parent(Data),
    append(Data, [a_i], NewData),
    depth_first_search(Parent, NewData).

% performing depth first search
depth_first_search(Parent, Data) :-
    find_next(Data, Next_Agent),
    (Next_Agent == none ->
        send_message(a_i, Parent, tree, Data)
    ;
    assign_child(Next_Agent),
    send_message(a_i, Next_Agent, tree, Data),
    wait_message(Next_Agent, a_i, tree, NewData),
    depth_first_search(Parent, NewData)
    ).
```

Intuitively, upon receiving a *tree* message from the agent *Parent* enclosed with data *Data*, the agent  $a_i$  will execute the clause `generate_tree(Parent, Data)`. Specifically:

- It executes the clause `assign_parent/1`, where it adds to its  $\Pi_{a_i}$  a fact of the form *parent(Parent)*;
- It executes the clause `assign_pseudo_parent/1` which adds to its  $\Pi_{a_i}$  facts of the form *pseudo\_parent(X)*, where *X* is a neighboring agent of  $a_i$  that appears in *Data* such that  $X \neq \textit{Parent}$ ;
- It adds itself (i.e., *a\_i*) to the list of visited agents;
- It starts performing a DFS, by executing the rule `depth_first_search/2`.

In order to perform a DFS, the agent  $a_i$  will execute the rule `find_next/2` to select a neighboring agent that will be visited next; this selection is based on some heuristics (i.e., the unvisited neighbor agent with the greatest number of neighbors). If such an agent *Next\_Agent* exists (i.e.,  $\textit{Next\_Agent} \neq \textit{none}$ ), then  $a_i$  will

- execute the rule `assign_child/1`, used to add to its  $\Pi_{a_i}$  a fact of the form `children(Next_Agent)`;
- send a `tree` message to the agent `Next_Agent`;
- wait for the reply message from the agent `Next_Agent`, which will provide the updated list `NewData` of visited agents;
- recursively execute the rule `depth_first_search/2`.

Otherwise, if there is no agent `Next_Agent` (i.e., `Next_Agent` is equal to `none`), then the agent  $a_i$  will reply a `tree` message to its agent `Parent`. This implies that the agent  $a_i$  has finished performing DFS at its branch.

When the agent  $a_i$  is chosen to be the root of the pseudo-tree, it executes the rule `generate_tree(master, [])` immediately without waiting for the `tree` message from other agents. We note that an agent whose parent agent is `master` will be the root of the pseudo-tree. It is also worth to notice that, at the end of this phase, the information about the parent, pseudo-parents, and child agents of each agent  $a_i$  are added to  $\Pi_{a_i}$  via facts of the forms `parent/1`, `pseudo_parent/1`, and `children/1`, respectively.

#### *Lemma 1*

In ASP-DPOP, Phase 1 requires a linear number of messages in  $n$  where  $n$  is the number of agents.

#### *Proof*

We first prove that Phase 1 terminates. In fact, each agent  $a_i$  in executing `depth_first_search/2` will perform the rule `find_next/2` to select a neighboring agent, i.e., `Next_Agent`, that is *not* in the set of visited agents to send a `tree` message to. `Next_Agent` can be seen as an unvisited neighboring agent of the agent  $a_i$ .

The agent  $a_i$  then waits to receive the `tree` message from `Next_Agent` enclosing an updated set of visited agents, and again send a `tree` message to another unvisited neighboring agent if it exists. We notice that the updated set of visited agents will be expanded with at least one agent that is `Next_Agent` since `Next_Agent` will add itself to the set of visited agents beyond receiving the `tree` message from the agent  $a_i$ .

Therefore, every agent  $a_i$  will send at most  $|N(a_i)|$  `tree` messages to its child agents, where  $N(a_i)$  is the set of the neighboring agents of the agent  $a_i$ . If there is no unvisited neighboring agent left, the agent  $a_i$  will send a `tree` message to its parent agent together with the most updated set of visited agents, and terminates executing `depth_first_search/2`. Thus, it terminates performing the clause `generate_tree(Parent, Data)`. As a consequence, we can conclude that Phase 1 terminates.

Furthermore, considering a pseudo-tree that is generated at the end of Phase 1. We can realize that the set of visited agents which are passing among agents is expanded with a non-root agent if and only if there is a `tree` message sent from a parent agent to its child agent downward the pseudo-tree. It is worth to remind that the agent who is selected to be the root of the pseudo-tree adds itself to the

set of visited agents at the beginning. Thus, there are  $n - 1$  *tree* messages that are sent downward the pseudo-tree. Moreover, every agent except the root agent will send exactly one *tree* message to its parent agent upward the pseudo-tree. Therefore, there are  $n - 1$  *tree* messages that are sent upward the pseudo-tree. Accordingly, in total there are  $2 \times (n - 1)$  *tree* messages produced in Phase 1. This proves Lemma 1.  $\square$

### 3.4.4 Computing the UTIL message: Phase 2

In the following, for simplicity, given an agent  $a_i$ , we assume that  $a_p = P_i$ . We will use  $a_p$  and  $P_i$  interchangeably. In this phase, each ASP-DPOP agent generates an ASP program for computing the UTIL message that will be sent to its parent. In more detail, each agent  $a_i$  executes the following clause:

```
% Phase 2: UTIL Propagation Phase
perform_Phase_2(ReceivedUTILMessages):-
    compute_separator(ReceivedUTILMessages, Separator),
    assert(separatorlist(Separator),
    compute_related_constraints(ConstraintList),
    assert(constraintlist(ConstraintList),
    generate_UTIL_ASP(Separator, ConstraintList),
    solve_answer_set1(ReceivedUTILMessages, Answer),
    store(Answer),
    send_message(a_i, a_p, util, Answer).
```

In particular, each agent  $a_i$  with a set of child agents  $C_i$ :

- Waits to receive all of the UTIL messages from its children and combines them into a set of ASP facts. Let  $M_{a_k}$  be the encoding of the UTIL message  $UTIL_{a_k}^{a_i}$ . We define a list *ReceivedUTILMessages* as follows.

$$ReceivedUTILMessages = \bigcup_{a \in C_i} M_a. \quad (3.33)$$

When  $a_i$  is a leaf ( $C_i = \emptyset$ ), we set *ReceivedUTILMessages* = [].

- Computes its separator  $sep_i$  by executing `compute_separator/2`. This is realized using (i) the information about its parent and pseudo-parent agents added in  $\Pi_{a_i}$  during Phase 1, and (ii) the information about ancestors of the agent  $a_i$  that are directly connected with descendants of the agent  $a_i$ , via facts of the form *table\_info*, contained in the UTIL messages received from its child agents;
- Computes the set  $R_{a_i}$  (*ConstraintList*) of the related constraints (i.e., executing `compute_related_constraints/1`) that is defined as (2.5).
- Generates the information for its UTIL message (i.e., executing `generate_UTIL_ASP/2`). Specifically, `generate_UTIL_ASP/2` first creates a logic program, denoted by  $I_{a_i}$ , from the separator list, the constraint list, and the information from  $M_{a_k}$  where  $a_k \in C_i$ . It then computes the answer set of  $\Pi_{a_i} \cup I_{a_i} \cup (\bigcup_{a \in C_i} M_a)$  that contains the encoded UTIL message of the agent  $a_i$ . Assume that

- $sep_i = \{x_{s_1}, \dots, x_{s_k}\}$  (i.e., *Separator* =  $[x_{s_1}, \dots, x_{s_k}]$  is the separator list of  $a_i$ );
- $R_{a_i} = \{f_{r_1}, \dots, f_{r_{k'}}\}$  and  $scp(f_{r_j}) = \{x_{r_{j_1}}, \dots, x_{r_{j_w}}\}$  for  $1 \leq j \leq k'$  (i.e., *ConstraintList* =  $[f_{r_1}, \dots, f_{r_{k'}}]$ );
- $C_i = \{a_{c_1}, \dots, a_{c_l}\}$  and each  $UTIL_{a_{c_t}}^{a_i}$  has  $x_{c_{t_1}}, \dots, x_{c_{t_z}}$  as its dimensions for  $1 \leq t \leq l$ ; and
- $a_p$  is the parent agent of the agent  $a_i$ .

generate\_UTIL\_ASP/2 creates  $I_{a_i}$  with the following rules:

$$\begin{aligned}
 \text{table\_row\_}a_i(U, X_{s_1}, \dots, X_{s_k}) \leftarrow & \\
 & f_{r_1}(V_{r_1}, X_{r_{1_1}}, \dots, X_{r_{1_w}}), \\
 & \dots \\
 & f_{r_{k'}}(V_{r_{k'}}, X_{r_{k'_1}}, \dots, X_{r_{k'_w}}), \\
 & V_{r_1} != \#inf, \dots, V_{r_{k'}} != \#inf, \\
 & \text{table\_max\_}a_{c_1}(U_{c_1}, X_{c_{1_1}}, \dots, X_{c_{1_z}}), \\
 & \dots \\
 & \text{table\_max\_}a_{c_l}(U_{c_l}, X_{c_{l_1}}, \dots, X_{c_{l_z}}), \\
 & U = V_{r_1} + \dots + V_{r_{k'}} + U_{c_1} + \dots + U_{c_l}.
 \end{aligned}
 \tag{3.34}$$

$$\begin{aligned}
 \text{table\_max\_}a_i(U, X_{s_1}, \dots, X_{s_k}) \leftarrow & \\
 & \text{value}(x_{s_1}, X_{s_1}), \\
 & \dots \\
 & \text{value}(x_{s_k}, X_{s_k}), \\
 & \text{table\_row\_}a_i(\neg, X_{s_1}, \dots, X_{s_k}), \\
 & U = \#max\{V : \text{table\_row\_}a_i(V, X_{s_1}, \dots, X_{s_k})\}.
 \end{aligned}
 \tag{3.35}$$

$$\begin{aligned}
 & \text{table\_info}(a_i, a_{s_1}, x_{s_1}, lb_{s_1}, ub_{s_1}). \\
 & \dots \\
 & \text{table\_info}(a_i, a_{s_k}, x_{s_k}, lb_{s_k}, ub_{s_k}).
 \end{aligned}
 \tag{3.36}$$

generate\_UTIL\_ASP/2 uses the information in  $UTIL_{a_{c_t}}^{a_i}$  and  $\Pi_{a_i}$  to generate the facts (3.36). In addition, each variable in the rules (3.34)–(3.36) corresponds to a variable name (e.g.,  $X_{s_1}$  corresponds to  $x_{s_1}$  in the separator list;  $X_{c_{1_1}}$  corresponds to  $x_{c_{1_1}}$  in dimensions of  $UTIL_{a_{c_1}}^{a_i}$ ; etc.). Therefore, due to the definition of the separator of  $a_i$  and  $sep_i = \{x_{s_1}, \dots, x_{s_k}\}$ ,  $X_{s_1}, \dots, X_{s_k}$  are guaranteed to occur on the right hand side of (3.34). In other words,  $I_{a_i}$  is a *safe* program.

Intuitively, the rule of the form (3.34) creates the joint table for  $a_i$ —that is similar to the result of flattening  $(\oplus_{a_{c_t} \in C_i} UTIL_{a_{c_t}}^{a_i}) \oplus (\oplus_{f \in R_{a_i}} f)$  into a table—given  $UTIL_{a_{c_t}}^{a_i}$  and  $R_{a_i}$ . In addition, (3.35) computes the optimal utilities for each value combination of variables in the separator list.

- Computes an answer set  $A$  of the program  $\Pi_{a_i} \cup I_{a_i} \cup \bigcup_{a \in C_i} M_a$  by executing solve\_answer\_set1/2, and extracts from  $A$  the information for the UTIL message (i.e., *Answer*) that will be sent from the agent  $a_i$  to the agent  $a_p$ .
- Asserts the information in *Answer* for later use in Phase 3 (i.e., executes store(*Answer*)).
- Sends encoded  $UTIL_{a_i}^{a_p}$  to the parent agent  $a_p$  (i.e., executes send\_message/4).

*Example 13*

As an example, we refer to the DCOP in Example 1. Specifically, we illustrate  $I_{a_2}$  generated by the agent  $a_2$ . *ReceivedUTILMessages* for the agent  $a_2$  is the set of ASP facts given in Example 10, *Separator* =  $[x_1]$ , and *ConstraintList* =  $[x_1\_cons\_x2]$ . The program  $I_{a_2}$  includes the following rules:

$$\begin{aligned} & \text{table\_row\_}a_2(U, X_1) \leftarrow x_1\_cons\_x2(V_0, X_1, X_2), \\ & \quad V_0 \neq \#inf, \\ & \quad \text{table\_max\_}a_3(V_1, X_1, X_2), \\ & \quad U = V_0 + V_1. \\ & \text{table\_max\_}a_2(U, X_1) \leftarrow \text{value}(x_1, X_1), \\ & \quad \text{table\_row\_}a_2(\_, X_1) \\ & \quad U = \#max\{V : \text{table\_row\_}a_2(V, X_1)\}. \\ & \text{table\_info}(a_2, a_1, x_1, 0, 1) \leftarrow \end{aligned}$$

The relationship between the ASP-based computation and Algorithm 1 is established in the following lemma.

*Lemma 2*

Let us consider a DCOP  $\mathcal{M}$ , an agent  $a_i \in \mathcal{A}$ , and a pseudo-tree  $T$ . Let  $a_i$  be an agent with  $C_i = \{a_{c_1}, \dots, a_{c_l}\}$  and  $M_{a_{c_t}}$  be the encoded  $UTIL_{a_{c_t}}^{a_i}$  for  $1 \leq t \leq l$ . Furthermore, let us assume that  $a_p$  is the parent of  $a_i$ ,  $sep_i = \{x_{s_1}, \dots, x_{s_k}\}$ , and  $R_{a_i} = \{f_{r_1}, \dots, f_{r_{k'}}\}$  and  $scp(f_{r_j}) = \{x_{r_{j_1}}, \dots, x_{r_{j_w}}\}$  for  $1 \leq j \leq k'$ . Then, the program  $\Pi_{a_i} \cup I_{a_i} \cup (\bigcup_{a \in C_i} M_a)$  has a unique answer set  $A$  and

- $\text{table\_row\_}a_i(u, v_{s_1}, \dots, v_{s_k}) \in A$  iff there exists a value combination  $X$  for variables in  $scp(JOIN_{a_i}^{P_i}(X))$  such that  $JOIN_{a_i}^{P_i}(X) = u$  where  $\{x_{s_1} = v_{s_1}, \dots, x_{s_k} = v_{s_k}\} \subseteq X$  and  $u \neq -\infty$ ; and
- $\text{table\_max\_}a_i(u, v_{s_1}, \dots, v_{s_k}) \in A$  iff  $UTIL_{a_i}^{P_i}(x_{s_1} = v_{s_1}, \dots, x_{s_k} = v_{s_k}) = u$  and  $u \neq -\infty$ .

*Proof*

Since  $I_{a_i}$  is safe and  $\Pi_{a_i} \cup I_{a_i} \cup (\bigcup_{a \in C_i} M_a)$  is a positive program, it has a unique answer set.

By the definition of answer set,  $\text{table\_row\_}a_i(u, v_{s_1}, \dots, v_{s_k}) \in A$  iff that there exists a rule  $r$  of the form (3.34) such that  $\text{table\_row\_}a_i(u, v_{s_1}, \dots, v_{s_k})$  is the head of  $r$ . It means that there exists a value assignment  $Y$  for the variables occurring in  $r$  such that the following conditions hold:

- $\{x_{s_1} = v_{s_1}, \dots, x_{s_k} = v_{s_k}\} \subseteq Y$ ;
- for each  $1 \leq j \leq k'$ , there exists  $v_{r_j} \neq \#inf$  such that  $f_{r_j}(v_{r_{j_1}}, \dots, v_{r_{j_w}}) = v_{r_j}$  and  $\{x_{r_{j_1}} = v_{r_{j_1}}, \dots, x_{r_{j_w}} = v_{r_{j_w}}\} \subseteq Y$ ; and
- for each  $1 \leq t \leq l$ , there exists  $u_{c_t}$  such that  $\text{table\_max\_}a_{c_t}(u_{c_t}, v_{c_{t_1}}, \dots, v_{c_{t_z}}) \in A$  and  $\{x_{c_{t_1}} = v_{c_{t_1}}, \dots, x_{c_{t_z}} = v_{c_{t_z}}\} \subseteq Y$ . By the construction of the algorithm,  $\text{table\_max\_}a_{c_t}(u_{c_t}, v_{c_{t_1}}, \dots, v_{c_{t_z}}) \in A$  implies that  $UTIL_{a_{c_t}}^{a_i}(x_{c_{t_1}} = v_{c_{t_1}}, \dots, x_{c_{t_z}} = v_{c_{t_z}}) = u_{c_t}$  and  $u_{c_t} \neq -\infty$ .

The conclusion of the first item follows directly from the definitions of the UTIL message and the  $\oplus$  operator (Definitions 3 and 4) and the above conditions.

The second item of the lemma follows from the first item, the condition  $V_{r_1} != \#inf, \dots, V_{r_{k'}} != \#inf$  in the rule (3.34), and Definition 5.  $\square$

Lemma 2 implies that Phase 2 of ASP-DPOP computes the same UTIL messages as DPOP, except that UTIL messages in ASP-DPOP omit the value assignments whose associated utilities are  $-\infty$ .

### 3.4.5 Computing the VALUE message: Phase 3

Each ASP-DPOP agent computes the optimal value for its variables and sends an encoded VALUE message to its children. The process is described by the following Prolog rule:

```
% Phase 3: VALUE Propagation Phase
perform_Phase_3(ReceivedVALUEMessage) :-
    separatorlist(Separator),
    constraintlist(ConstraintList),
    generate_VALUE_ASP(Separator,ConstraintList),
    solve_answer_set2(ReceivedVALUEMessage, Answer),
    send_message_to_children(a_i, value, Answer).
```

In particular, the agent  $a_i$

- waits to receive the encoded VALUE message, denoted by  $M'_p$ , from its parent agent  $P_i$ . If the agent  $a_i$  does not have a parent, i.e., it is the root of the tree, we set *ReceivedVALUE Message* = [];
- retrieves  $sep_i$  (i.e., *Separator*) computed in Phase 2;
- retrieves  $R_{a_i}$  (i.e., *ConstraintList*) computed in Phase 2;
- executes the rule *generate\_VALUE\_ASP/2* to create an ASP program, denoted by  $I'_{a_i}$ , from the separator list, the constraint list, and the information from  $M_{a_k}$  where  $a_k \in C_i$  from Phase 2. Assume that
  - $sep_i = \{x_{s_1}, \dots, x_{s_k}\}$  (i.e., *Separator* =  $[x_{s_1}, \dots, x_{s_k}]$  is the separator list of  $a_i$ );
  - $R_{a_i} = \{f_{r_1}, \dots, f_{r_{k'}}\}$  and  $scp(f_{r_j}) = \{x_{r_{j_1}}, \dots, x_{r_{j_w}}\}$  for  $1 \leq j \leq k'$  (i.e., *ConstraintList* =  $[f_{r_1}, \dots, f_{r_{k'}}]$ );
  - $C_i = \{a_{c_1}, \dots, a_{c_l}\}$  and each  $UTIL_{a_{c_t}}^{a_i}$  has  $x_{c_{t_1}}, \dots, x_{c_{t_z}}$  as its dimensions for  $1 \leq t \leq l$ ; and
  - the set of variables owned by the agent  $a_i$  is  $\alpha_i = \{x_{i_1}, \dots, x_{i_q}\}$ .

generate\_VALUE\_ASP/2 creates the logic program  $I'_{a_i}$  with following rules:

$$\begin{aligned}
 \{\text{row}(U, X_{i_1}, \dots, X_{i_q})\} &\leftarrow \text{solution}(\alpha(x_{s_1}), x_{s_1}, X_{s_1}), \\
 &\dots \\
 &\text{solution}(\alpha(x_{s_k}), x_{s_k}, X_{s_k}), \\
 &\text{table\_max\_}a_i(U, X_{s_1}, \dots, X_{s_k}), \\
 &f_{r_1}(V_{r_1}, X_{r_{1_1}}, \dots, X_{r_{1_w}}), \\
 &\dots \\
 &f_{r_{k'}}(V_{r_{k'}}, X_{r_{k'_1}}, \dots, X_{r_{k'_w}}), \\
 &\text{table\_max\_}a_{c_1}(U_{c_1}, X_{c_{1_1}}, \dots, X_{c_{1_z}}), \\
 &\dots \\
 &\text{table\_max\_}a_{c_l}(U_{c_l}, X_{c_{l_1}}, \dots, X_{c_{l_z}}), \\
 U &== V_{r_1} + \dots + V_{r_{k'}} + U_{c_1} + \dots + U_{c_l}.
 \end{aligned} \tag{3.37}$$

$$\leftarrow \text{not } 1\{\text{row}(U, X_{i_1}, \dots, X_{i_q})\}1. \tag{3.38}$$

$$\begin{aligned}
 \text{solution}(a_i, x_{i_1}, X_{i_1}) &\leftarrow \text{row}(U, X_{i_1}, \dots, X_{i_q}), \\
 &\dots \\
 \text{solution}(a_i, x_{i_q}, X_{i_q}) &\leftarrow \text{row}(U, X_{i_1}, \dots, X_{i_q}).
 \end{aligned} \tag{3.39}$$

Intuitively, the rule of the form (3.37) and the constraint of the form (3.38) select an optimal row based on: (i) The computation as done in Phase 2 (i.e., using the facts of the form `table_max_` $a_i$  that are stored in Phase 2), and (ii) (for non-root agent only) the VALUE message that is received from its parent (i.e., facts of the form `solution/3`). The selected optimal row will define the optimal value of the variables using rules of the form (3.39). Similar argument for the safety of  $I_{a_i}$  allows us to conclude that  $I'_{a_i}$  is also a *safe* program.

- Executes `solve_answer_set2/2`, that executes the ASP solver to compute an answer set of the program  $\Pi_{a_i} \cup I'_{a_i} \cup M'_{P_i} \cup (\bigcup_{a \in C_i} M_a)$ . From that answer set, it collects all facts of the form `solution(a, x, v)` and returns them as *Answer*—i.e., the encoding of the VALUE message from the agent  $a_i$  to its child agents.
- Executes `send_message_to_children/3` where it sends *value* message with *Answer* as its data to each child agent (i.e., executing the respected clauses `send_message/4` multiple times).

#### Example 14

As an example, we refer to the DCOP in Example 1 to illustrate  $I'_{a_2}$  generated by the agent  $a_2$ . *ReceivedUTILMessages* for the agent  $a_2$  is the set of ASP facts given in Example 10, *Separator* =  $[x_1]$ , *ConstraintList* =  $[x1\_cons\_x2]$ , and  $\alpha_2 = \{x_2\}$ . The



program  $I'_{a_2}$  includes following rules:

$$\begin{aligned} \{\text{row}(U, X_2)\} &\leftarrow \text{solution}(a_1, x_1, X_1), \\ &\quad \text{table\_max\_}a_2(U, X_1), \\ &\quad x_1\text{-cons\_}x_2(U_0, X_1, X_2), \\ &\quad \text{table\_max\_}a_3(U_1, X_1, X_2) \\ &\quad U == U_0 + U_1. \\ &\leftarrow \text{not } 1\{\text{row}(U, X_2)\}1. \\ \text{solution}(a_2, x_2, X_2) &\leftarrow \text{row}(U, X_2). \end{aligned}$$

*Lemma 3*

Let us consider a DCOP  $\mathcal{M}$ , and an agent  $a_i \in \mathcal{A}$  in a pseudo-tree  $T$ . Let  $a_i$  be an agent with  $C_i = \{a_{c_1}, \dots, a_{c_l}\}$  and  $M_{a_{c_t}}$  be the encoding of  $UTIL_{a_{c_t}}^{a_i}$  for  $1 \leq t \leq l$ . Furthermore, assume that  $P_i$  is the parent agent of the agent  $a_i$ ,  $sep_i = \{x_{s_1}, \dots, x_{s_k}\}$ ,  $R_{a_i} = \{f_{r_1}, \dots, f_{r_{k'}}\}$  where  $scp(f_{r_j}) = \{x_{r_{j_1}}, \dots, x_{r_{j_w}}\}$  for  $1 \leq j \leq k'$ ,  $\alpha_i = \{x_{i_1}, \dots, x_{i_q}\}$ , and  $M_{P_i}$  encodes  $VALUE_{P_i}^{a_i}$ . Let  $Q = \Pi_{a_i} \cup I'_{a_i} \cup M'_{P_i} \cup (\bigcup_{a \in C_i} M_a)$ . Then,

- for each answer set  $A$  of  $Q$ , the assignment  $x_{i_j} = v_{i_j}$ , where  $\text{solution}(a_i, x_{i_j}, v_{i_j}) \in A$  for  $1 \leq j \leq q$  belongs to a solution of  $\mathcal{M}$ ; and
- if  $x_{i_1} = v_{i_1}, \dots, x_{i_q} = v_{i_q}$  is a value assignment for variables in  $\alpha_i$  that belongs to a solution of  $\mathcal{M}$ , which contains  $VALUE_{P_i}^{a_i}$ , then  $Q$  has an answer set  $A$  containing  $\{\text{solution}(a_i, x_{i_j}, v_{i_j}) \mid 1 \leq j \leq q\} \cup M'_i$ .

*Proof*

Based on the construction of  $I'_{a_i}$ , it is possible to see that there exists at least one rule of the form (3.37) in  $Q$ . Observe that if the agent  $a_i$  is the root of the pseudo-tree  $T$ . Then,  $M'_{P_i} = \emptyset$  and the rule (3.37) does not contain the atom of the form  $\text{solution}(a, x, v)$ . Since the program is safe and positive, we have that  $Q$  is consistent.

Because of the rule (3.38), each answer set  $A$  of  $Q$  contains exactly one atom of the form  $\text{row}(u, v_{i_1}, \dots, v_{i_q})$ . Also, from the rule (3.37), we have that if  $\text{row}(u, v_{i_1}, \dots, v_{i_q}) \in A$ , then there exists some  $\text{table\_max\_}a_i(u, v_{s_1}, \dots, v_{s_k}) \in A$ , which indicates that  $u$  is the optimal utility corresponding to the assignment  $x_{s_i} = v_{s_i}$  for  $1 \leq i \leq k$  (Lemma 2). From the correctness of DPOP, this means that  $\text{row}(u, v_{i_1}, \dots, v_{i_q})$  encodes an optimal value assignment for variables owned by the agent  $a_i$ . This proves the first item.

Assume that  $x_{i_1} = v_{i_1}, \dots, x_{i_q} = v_{i_q}$  is a value assignment for variables in  $\alpha_i$  that belongs to a solution of  $\mathcal{M}$ , which contains  $VALUE_{P_i}^{a_i}$ . Then, by the completeness of DPOP and Lemma 2, this implies that there exists some  $\text{table\_max\_}a_i(u, v_{s_1}, \dots, v_{s_k})$  such that  $VALUE_{P_i}^{a_i}$  contains  $x_{s_i} = v_{s_i}$  for  $1 \leq i \leq k$ . As such, there must exist the values for  $f_{r_j}(\cdot)$  and  $\text{table\_max\_}a_{c_t}(\cdot)$  such that there is a rule of the form (3.37) whose head is  $\text{row}(u, v_{i_1}, \dots, v_{i_q})$ . This means that  $Q$  has an answer set containing  $\text{row}(u, v_{i_1}, \dots, v_{i_q})$ , which proves the second item of the lemma. □

## 3.4.6 ASP-DPOP

In this section, we will show the clause for ASP-DPOP agents to perform Phase 1, Phase 2, and Phase 3 consecutively. For simplicity, we omit the fragment of code of ASP-DPOP agents that allow them to cooperate with each other to select one agent as the root of the pseudo-tree—since it depends on the scores that are assigned to agents, according to a heuristic function.

Let us remind that if an agent  $a_k$  is the root of the pseudo-tree, a fact of the form  $\text{parent}(\text{master}, a_k)$  will be added to  $\Pi_{a_k}$ . After an agent is selected as the root of the pseudo-tree, each agent will execute the clause `start`. Considering an agent  $a_i$ , the clause `start` of the agent  $a_i$  is described as follows:

```
% Perform Phase 1, Phase 2, and Phase 3
start:-
    (parent(master, a_i) ->
        generate_tree(master, [])
        ;
        wait_message(Parent, a_i, tree, Data),
        generate_tree(Parent, Data)
    ),
    (children(_) ->
        get_UTILMessages_from_all_children(ReceivedUTILMessages),
        perform_Phase_2(ReceivedUTILMessages)
        ;
        perform_Phase_2([])
    ),
    (parent(master, a_i) ->
        perform_Phase_3([])
        ;
        wait_message(Parent, a_i, value, ReceivedVALUEMessage),
        perform_Phase_3(ReceivedVALUEMessage)
    ).
```

In particular, each agent  $a_i$ :

- Checks whether  $a_i$  is the root of the pseudo-tree; this is realized by checking whether the fact of the form  $\text{parent}(\text{master}, a_i)$  is in  $\Pi_{a_i}$ :
  - If the agent  $a_i$  is the root of the pseudo-tree, it will execute `generate_tree(master, [])` that is defined in Section 3.4.3; otherwise,
  - If the agent  $a_i$  is not the root of the pseudo-tree, it will execute `wait_message(Parent, a_i, tree, Data)`. Upon receiving the `tree` message from the agent `Parent` who is later assigned as its parent agent, the agent  $a_i$  will execute `generate_tree(Parent, Data)` that is defined in Section 3.4.3.
- Checks whether the agent  $a_i$  is (resp. is not) a leaf of the pseudo-tree (i.e., this is realized by checking whether the fact of the form `children/1` is not (resp. is) in  $\Pi_{a_i}$ , respectively):
  - If the agent  $a_i$  is not a leaf of the pseudo-tree, it executes `get_UTILMessages_from_all_children(ReceivedUTILMessages)`. Intuitively, the clause `get_UTILMessages_from_all_children/1` iteratively executes `wait_message(From, a_i, util, Data)` until the agent  $a_i$

receives all *util* messages from all of its child agents. The contents (i.e., *Data*) of all *util* messages are combined into *ReceivedUTILMessages*. Then, the agent  $a_i$  executes *perform\_Phase\_2*(*ReceivedUTILMessages*) that is defined in Section 3.4.4; otherwise,

- If the agent  $a_i$  is a leaf of the pseudo-tree, it executes *perform\_Phase\_2*( $[]$ ) that is defined in Section 3.4.4.
- Checks whether the agent  $a_i$  is the root of the pseudo-tree or not:
  - If the agent  $a_i$  is the root of the pseudo-tree, it executes *perform\_Phase\_3*( $[]$ ) that is defined in Section 3.4.5; otherwise,
  - If the agent  $a_i$  is not the root of the pseudo-tree, it executes *wait\_message*(*Parent*, *a.i*, *value*, *ReceivedVALUEMessage*) to wait for *value* message from its parent agent. Then, the agent  $a_i$  executes *perform\_Phase\_3*(*ReceivedVALUEMessage*) that is defined in Section 3.4.5.

## 4 Theoretical analysis

In this section, we present some theoretical properties of ASP-DPOP including its soundness, completeness, and complexity.

### 4.1 Soundness and completeness

The soundness and completeness of ASP-DPOP follow from Lemmas (2) and (3) and the soundness and completeness of DPOP.

#### *Proposition 1*

ASP-DPOP is sound and complete in solving DCOPs.

#### *Proof*

Let us summarize how ASP-DPOP solves a DCOP  $\mathcal{M}$ :

- In Phase 1, each ASP-DPOP agent runs distributed DFS to generate a pseudo-tree. At the end of this phase, the information about the parent, pseudo-parents, and child agents of each agent  $a_i$  are added to  $\Pi_{a_i}$  via facts of the forms *parent/1*, *pseudo\_parent/1* and *children/1*, respectively;
- In Phase 2, each ASP-DPOP agent  $a_i$ : (i) waits to receive the encoding of UTIL messages from all of its child agents (for non-leaf agents only), and (ii) generates the ASP program  $I_{a_i}$  to compute its encoded UTIL message as an answer set of  $\Pi_{a_i} \cup I_{a_i} \cup \text{ReceivedUTILMessages}$ ;
- In Phase 3, each ASP-DPOP agent  $a_i$ : (i) waits to receive the encoded VALUE message from its parent agent (for non-root agent only), and (ii) generates the ASP program  $I'_{a_i}$  to compute its encoded VALUE message as an answer set of  $\Pi_{a_i} \cup I'_{a_i} \cup M'_{P_i} \cup \text{ReceivedUTILMessages}$ ;

The soundness and completeness of ASP-DPOP follow from the soundness and completeness of DPOP and the following observations:

- Phase 1 of ASP-DPOP generates a possible pseudo-tree of  $\mathcal{M}$ .
- Assuming that ASP-DPOP and DPOP use the same pseudo-tree  $T$ , then
  - Phase 2 of ASP-DPOP computes the same UTIL messages as DPOP except that they omit the value assignments whose associated utilities are  $-\infty$  (Lemma 2). However, for DPOP, ignoring those value assignments in UTIL message will not affect DCOP solution since such value assignments are not included in any solution (i.e., otherwise the total utility is  $-\infty$ ).
  - Phase 3 of ASP-DPOP computes all possible solutions (VALUE messages) as DPOP (Lemma 3).

□

## 4.2 Complexity

Given  $d = \max_{1 \leq i \leq n} |D_i|$  and  $w^* = \max_{1 \leq i \leq n} |sep_i|$ , where  $n$  is the number of agents,<sup>7</sup> we have the following properties:

### Property 1

The number of messages required by ASP-DPOP is bounded by  $\mathcal{O}(n)$ .

### Proof

In ASP-DPOP, one can observe that: (1) Phase 1 requires a linear number of messages in  $n$  (Lemma 1); (2) Phase 2 requires  $(n - 1)$  UTIL messages; and (3) Phase 3 requires  $(n - 1)$  VALUE messages. Thus, the number of messages required by ASP-DPOP is bounded by  $\mathcal{O}(n)$ . □

### Property 2

The size of messages required by ASP-DPOP is bounded by  $\mathcal{O}(d^{w^*})$ .

### Proof

In ASP-DPOP,

- Phase 1 produces messages whose size is linear in  $n$ . This is because the *tree* message is of the form *send\_message(a.i, Next\_Agent, tree, Data)* where the content (*Data*) that dominates the size of the *tree* message is the set of visited agents whose size is linear in  $n$ ;
- Phase 2 produces encoded UTIL messages; each message consists of: (i) a fact of the form *table\_max\_a<sub>i</sub>* for each cell in the corresponding UTIL message in DPOP where its associated optimal utility is not  $-\infty$ , and (ii)  $|sep_i|$  facts of the form *table\_info*. Therefore, the size of encoded UTIL messages is bounded by  $\mathcal{O}(d^{w^*})$  as the bounded size of UTIL messages in DPOP (Petcu and Faltings 2005a); and
- Phase 3 produces encoded VALUE messages; each message consists of a fact of the form *solution/3* for each value assignment of a variable in the corresponding VALUE message in DPOP. Therefore, Phase 3 produces encoded VALUE messages whose sizes are bounded by  $\mathcal{O}(|\mathcal{X}|) = \mathcal{O}(n)$  as we assume each agent owns exactly one variable.

<sup>7</sup>  $w^*$  is also known as the *induced width* of a pseudo-tree (Dechter 2003).

Thus, the size of messages required by ASP-DPOP is bounded by  $\mathcal{O}(d^{w^*})$   $\square$

### Property 3

The memory requirements in ASP-DPOP are exponential and bounded by  $\mathcal{O}(d^{w^*})$ .

### Proof

In ASP-DPOP:

- In Phase 1, the memory requirements are bounded by  $\mathcal{O}(n)$  because it needs to keep track of the set of visited agents and the set of its neighboring agents;
- In Phase 2, the memory requirements are bounded by  $\mathcal{O}(d^{w^*})$  since, in computing the answer set of  $P = \Pi_{a_i} \cup I_{a_i} \cup \text{ReceivedUTILMessages}$ , the ASP solver needs to ground all the rules of the forms (3.34) and (3.35), and these dominate the number of other facts or ground instances of other rules in  $P$ ;
- In Phase 3, the memory requirements are bounded by  $\mathcal{O}(d^{w^*})$  since, in computing the answer set of  $P' = \Pi_{a_i} \cup I'_{a_i} \cup M'_{P_i} \cup \text{ReceivedUTILMessages}$ , the ASP solver needs to ground all rules of the form (3.37), where the number of facts of the form `table_max_a_i` is bounded by  $\mathcal{O}(d^{w^*})$  (see Property 2). Moreover, the number of such ground instances dominates the number of other facts and ground instances of other rules in  $P'$ .

Thus, the memory requirement in ASP-DPOP is exponential and bounded by  $\mathcal{O}(d^{w^*})$ .  $\square$

## 5 Experimental results

The goal of this section is to provide an experimental evaluation of ASP-DPOP. In particular, we compare ASP-DPOP against the original DPOP as well as other three implementations of complete DCOP solvers: *AFB*, *Hard Constraint-DPOP* (*H-DPOP*), and *Open-DPOP* (*ODPOP*). *AFB* (Gershman et al. 2009) is a complete search-based algorithm to solve DCOPs. *H-DPOP* (Kumar et al. 2008) is a complete DCOP solver that, in addition, propagates hard constraints to prune the search space. *ODPOP* (Petcu and Faltings 2006) is an optimization algorithm for DCOPs, which combines some advantages of both search-based algorithms and dynamic-programming-based algorithms. For completeness of the paper, in this section, we will first provide some background about these three solvers, discuss about FRODO platform (Léauté et al. 2009)—a publicly available implementation of DPOP, *AFB*, and *ODPOP*—and then provide some experimental results.

### 5.1 Background on *AFB*

The *AFB* algorithm (Gershman et al. 2009), to the best of our knowledge, is the most recent complete search-based algorithm to solve DCOPs. *AFB* makes use of a *Branch and Bound* scheme to identify a complete value assignment that minimizes the aggregate utility of all constraints. To do so, agents expand a partial value assignment as long as the lower bound on its aggregate utility does not exceed the

global bound, which is the aggregate utility of the best complete value assignment found so far.

In AFB, the state of the search process is represented by a data structure called *Current Partial Assignment (CPA)*. The CPA starts empty at some initializing agent, which records the value assignment of its own variable and sends the CPA to the next agent. The aggregate utility of a CPA is the accumulated utility of constraints involving the value assignment it contains. Each agent, upon receiving a CPA, adds a value assignment of its own variable such that the CPA's aggregate utility will not exceed the global upper bound. If it cannot find such an assignment, it backtracks by sending the CPA to the last assigning agent, requesting that agent to revise its value assignment.

Agents in AFB process and communicate CPAs asynchronously. An agent that succeeds to expand the value assignment of the received CPA sends forward copies of the updated CPA, requesting all unassigned agents to compute lower bound estimates of the aggregate utility of the current CPA. The assigning agent will receive these estimates asynchronously over time, and use them to update the lower bound of the CPA. Using these bounds, the assigning agent can detect if any expansion of this partial value assignment in the current CPA will cause it to exceed the global upper bound, and in such cases it will backtrack. Additionally, a time stamp mechanism for forward checking is used by agents to determine the most updated CPA and to discard obsolete CPAs.

## 5.2 Background on H-DPOP

In H-DPOP (Kumar *et al.* 2008), the authors consider how to leverage the hard constraints that may exist in the problem in a dynamic programming framework, so that only feasible partial assignments are computed, transmitted, and stored (Kumar *et al.* 2008). To this end, they encode combinations of assignments using *Constrained Decision Diagrams (CDDs)*. Basically, CDDs eliminate all inconsistent assignments and only include utilities corresponding to value combinations that are consistent. The resulting algorithm, H-DPOP, a hybrid algorithm that is based on DPOP, uses CDDs to rule out infeasible assignments, and thus compactly represents UTIL messages.

A CDD  $\mathcal{G} = (\Gamma, G)$  encodes the consistent assignments for a set of constraints  $\Gamma$  in a rooted direct acyclic graph  $G = (V, E)$  by means of constraint propagation. A node in  $G$  is called a CDD node. The terminal nodes are either *true* or *false* implying consistent or inconsistent assignment, respectively. By default, a CDD represents consistent assignments omitting the *false* terminal.

The H-DPOP algorithm leverages the pruning power of hard constraints by using CDDs to effectively reduce the message size. As in DPOP, H-DPOP has three phases: the pseudo-tree construction, the bottom-up UTIL propagation, and top-down VALUE propagation. The pseudo-tree construction and VALUE propagation phases are identical to ones of DPOP. In the UTIL propagation phase, the UTIL message, instead of being a multidimensional matrix, is a *CDDMessage*.

*Definition 6*

A CDDMessage  $M_i^j$  sent by an agent  $a_i$  to agent  $a_j$  is a tuple  $\langle \vec{u}, \mathcal{G} \rangle$  where  $\vec{u}$  is a vector of utilities, and  $\mathcal{G}$  is a CDD defined over variables in  $sep_i$ . The set of constraints for  $\mathcal{G}$  is  $\Gamma = \{f_j \mid scp(f_j) \subseteq sep_i\}$ .

In the UTIL propagation phase, H-DPOP defines different JOIN and PROJECTION operations. Observe that, based on Definition 6, an H-DPOP agent  $a_i$  can access a constraint whose scope is a subset of its separator, but that is not owned by  $a_i$  itself. For example, considering the DCOP in Example 1, in H-DPOP, the UTIL message sent from the agent  $a_3$  to the agent  $a_2$  will have information about the constraint  $x1\_cons\_x2$  that is not owned by the agent  $a_3$  since  $scp(x1\_cons\_x2) = \{x_1, x_2\} \subseteq sep_3$ . This might be undesirable in situations where the distribution of the computation is tied to privacy of information.

### 5.3 Background on ODPOP

ODPOP (Petcu and Faltings 2006) is an optimization algorithm for DCOP, which combines some advantages of both search-based algorithms and dynamic-programming-based algorithms. ODPOP always uses linear size messages, which is similar to search, and typically generates as few messages as DPOP. It does not always incur the worst complexity, which is the same with the complexity of DPOP, and on average it saves a significant amount of computation and information exchange. This is achieved because agents in ODPOP use a *best-first order* for value exploration and an *optimality criterion* that allows them to prove optimality without exploring all value assignments for variables in their separator. ODPOP also has three phases as DPOP:

**Phase 1 (DFS traversal)** is the same with Phase 1 in DPOP to construct a pseudo-tree.

**Phase 2 (ASK/GOOD)** is an iterative, bottom-up utility propagation process where each agent repeatedly sends ASK messages to its child agents, asking for valuations (GOODs), until it is able to compute the *suggested optimal value assignment* (GOOD) for variables in its separator. It then sends that GOOD, together with the respective utility that is obtained in the subtree rooted at this agent, as a GOOD message to its parent agent. This phase finishes when the root received enough GOODs to determine the optimal value assignment for its variables.

In more detail, in Phase 2, any child agent delivers to its parent agent a sequence of GOOD messages, each of which explores a different value assignment for variables in its separator, together with the corresponding utility. In addition, ODPOP uses a method to propagate GOODs so that every agent always reports its GOODs in order of *non-increasing utility*, provided that all of their child agents also follow this order. We can see that DPOP agents receive all GOODs that are grouped in single messages (i.e., UTIL messages). In contrast, ODPOP agents send GOODs on demand (i.e., when it receives an ASK message) individually and asynchronously as long as GOODs have non-increasing utilities.

As a consequence, each ODPOP agent  $a_i$  can determine when it has received *enough* GOODs from its child agents in order to be able to determine a GOOD

to send to its parent agent  $P_i$ . At that time,  $a_i$  will not send ASK message any more since any additional received GOODs will not affect the GOOD that was determined. If  $a_i$  later receives more ASK message from  $P_i$  for having next GOOD,  $a_i$  will continue to request more GOODs from its child agents until it can determine the next GOOD to report to  $P_i$ .

Since GOODs are always reported in order of non-increasing utility, the first GOOD that is generated at the root agent is the optimal value assignment for its variable. As a consequence, the root agent will be able to generate this solution without having to consider all value assignments for its variables.

**Phase 3 (VALUE propagation)** is similar to Phase 3 in DPOP. The root agent initiates the top-down VALUE propagation by sending a VALUE message to its child agents, informing them about its optimal value assignment for its variables. Subsequently, each agent  $a_i$ , upon receiving a VALUE message, will determine its optimal value assignment for its variables based on the computation (in Phase 2) of the first GOOD message generated whose associated value assignment is consistent with the one in the received VALUE message.

#### 5.4 Discussion on FRODO platform

In our experiment, we will compare the performance of ASP-DPOP against DPOP, AFB, and ODPOP; in particular, we use the implementation of the latter three systems that is publicly available in the FRODO platform (Léauté *et al.* 2009). It is important to observe that, at the implementation level, all DCOP solvers that are implemented within FRODO follow the *simplifying assumption* that each agent owns exactly one variable. This assumption is common practice in the DCOP literature (Modi *et al.* 2005; Petcu and Faltings 2005a; Gershman *et al.* 2009; Ottens *et al.* 2012). However, agents in DCOP problems used in our experiments own multiple variables. We will discuss in this subsection the pre-processing technique (i.e., *decomposition*, also known as *virtual agents*) that FRODO uses to transform a general DCOP with multiple variables per agent into a new DCOP with one variable per agent.

FRODO creates a virtual agent for each variable in a DCOP. A distinct variable is assigned to each virtual agent, so that this formulation satisfies the simplifying assumption. We say that a virtual agent  $a'_i$  belongs to a real agent  $a_i$  in DCOP if the virtual agent  $a'_i$  owns a variable that is owned by the real agent  $a_i$ . In FRODO, the solving algorithm is executed on each virtual agent, while intra-agent messages (i.e., messages are passed between virtual agents that belong to the same real agent) are only simulated and discounted in the calculation of computation cost (e.g., number of messages and messages' size).

Let  $M$  be a DCOP with multiple variables per agent, and  $M'$  be a new DCOP with one variable per agent that is constructed from  $M$ . Let us assume that we apply DPOP to solve both  $M$  and  $M'$ , using the same heuristics to construct the pseudo-trees. We can observe that each node in the pseudo-tree used to solve  $M'$  represents a virtual agent, while each node in a pseudo-tree of ASP-DPOP represents a real agent. It is possible to see that the number of inter-agent messages (i.e., messages



that are passed between virtual agents that belong to different real agents) produced in solving  $M'$  may be greater than the number of UTIL messages produced in solving  $M$ , depending on their respective pseudo-trees. Therefore, to minimize the total number of inter-agent messages, FRODO constructs pseudo-trees where virtual agents that belong to the same real agent stay as close as possible to each other.

It is important to summarize that, to handle a general DCOP with multiple variables per agent, FRODO first transforms it into a new DCOP with one variable per agent (introducing virtual agents), and then executes the resolution algorithms on each agent of the new DCOP. To the best of our knowledge, there is not any formal discussion about the relationship between pseudo-trees whose nodes represent real agents and pseudo-trees whose nodes represent virtual agents. However, it is believed that given a pseudo-tree  $T'$  whose nodes represent virtual agents, there always exists a pseudo-tree  $T$  whose nodes represent real agents such that  $T$  is compatible with  $T'$ . Intuitively, by compatible we mean that it is possible to construct  $T$  from  $T'$  as follows:

- If the root of  $T'$  is a node representing the virtual agent  $a'_i$  that belongs to a real agent  $a_i$ , the root of  $T$  is the node representing  $a_i$ ; and
- If there is at least one tree edge (resp. back edge) connecting two nodes that represent virtual agents  $a'_{i_1}$  and  $a'_{i_2}$  in  $T'$ , there is a tree edge (resp. back edge) connecting the two nodes that represent real agents  $a_{i_1}$  and  $a_{i_2}$  in  $T$  such that  $a'_{i_1}$  and  $a'_{i_2}$  belong to  $a_{i_1}$  and  $a_{i_2}$ , respectively.

It is worth to notice that, in our experiments, we ensure that all algorithms use the same heuristics to construct their pseudo-trees. We also observe that all pseudo-trees that are constructed using ASP-DPOP are compatible with the corresponding pseudo-trees that are constructed using FRODO.

### 5.5 Experimental results

We implement two versions of the ASP-DPOP algorithm—one that uses ground programs, which we call “ASP-DPOP (facts),” and one that uses non-ground programs, which we call “ASP-DPOP (rules).” In addition, as the observation made about H-DPOP in Section 5.2, we also implemented a variant of H-DPOP, called PH-DPOP, which stands for *Privacy-based H-DPOP* that restricts the amount of information that each agent can access to the amount common in most DCOP algorithms, including DPOP and ASP-DPOP. Specifically, agents in PH-DPOP can only access their own constraints and, unlike H-DPOP, cannot access their neighboring agents’ constraints.

In our experiments, we compare both versions of ASP-DPOP against DPOP (Petcu and Faltings 2005a), H-DPOP (Kumar et al. 2008), PH-DPOP, AFB (Gershman et al. 2009), and ODPOP (Petcu and Faltings 2006). We use a publicly-available implementation of DPOP, AFB, and ODPOP (Léauté et al. 2009) and an implementation of H-DPOP provided by the authors. We ensure that all algorithms use the same heuristics to construct their pseudo-trees for fair comparison. We measure the runtime of the algorithms using the simulated runtime metric (Sultanik

*et al.* 2007). All experiments are performed on a Quadcore 3.4 GHz machine with 16GB of memory. If an algorithm fails to solve a problem, it is due to memory limitations; other types of failures are specifically stated. We conduct our experiments on random graphs (Erdős and Rényi 1959), where we systematically modify the domain-independent parameters, and on comprehensive optimization problems in power networks (Gupta *et al.* 2013).

*Random graphs:* We create an  $n$ -node network, whose constraint density  $p_1$  produces  $\lfloor n \cdot (n - 1) \cdot p_1 \rfloor$  edges in total (Erdős and Rényi 1959). In our experiments, we vary the number of variables  $|\mathcal{X}|$ , the domain size  $|D_i|$ , the constraint density  $p_1$ , and the constraint tightness  $p_2$ . For each experiment, we vary only one parameter and fix the rest to their “default” values:  $|\mathcal{A}| = 5$ ,  $|\mathcal{X}| = 15$ ,  $|D_i| = 6$ ,  $p_1 = 0.6$ ,  $p_2 = 0.6$ . The timeout is set to  $5 \times 10^6$  ms. Table 2 shows the percentage of instances solved (out of 50 instances) and the average simulated runtime (in ms) for the solved instances. We do not show the results for ASP-DPOP (rules), as the utilities in the utility table are randomly generated, leading to no differences w.r.t. ASP-DPOP (facts). We make the following observations:

- ASP-DPOP is able to solve more problems than DPOP and is faster than DPOP when the problem becomes more complex (i.e., increasing  $|\mathcal{X}|$ ,  $|D_i|$ ,  $p_1$ , or  $p_2$ ). The reason is that ASP-DPOP is able to prune a significant portion of the search space thanks to hard constraints. ASP-DPOP does not need to explicitly represent the rows in the UTIL table that are infeasible, unlike DPOP. The size of the search space pruned increases as the complexity of the instance grows, resulting in a larger difference between the runtime of DPOP and ASP-DPOP.
- H-DPOP is able to solve more problems and solve them faster than DPOP, PH-DPOP, and ASP-DPOP. The reason is that agents in H-DPOP utilize more information about the neighbors’ constraints to prune values. In contrast, agents in ASP-DPOP and PH-DPOP only utilize information about their own constraints to prune values and agents in DPOP do not prune any values.
- ASP-DPOP is able to solve more problems and solve them faster than PH-DPOP. The reason is that agents in PH-DPOP, like agents in H-DPOP, use constraint decision diagram (CDD) to represent their utility tables, and it is expensive to maintain and perform join and project operations on this data structure. In contrast, agents in ASP-DPOP are able to capitalize on highly efficient ASP solvers to maintain and perform operations on efficient data structures thanks to their highly optimized grounding techniques and use of portfolios of heuristics.
- AFB is able to solve more problems and solve them faster than every other algorithm. We attribute this observation mainly to the relatively small number of variables in this experiment—i.e., the maximum number of variables in this experiment is 25 (see the first table in Table 2).
- ASP-DPOP is able to solve more problems and solve them faster than ODPOP when the problem becomes more complex (i.e., increasing  $|\mathcal{X}|$ ,  $|D_i|$ ,  $p_1$ , or  $p_2$ ). The reason is that ODPOP does not prune the search space based on hard

Table 2. Experimental results on random graphs

| $ \mathcal{X} $ | DPOP   |         | H-DPOP |      | PH-DPOP |         | AFB    |      | ODPOP  |         | ASP-DPOP |       |
|-----------------|--------|---------|--------|------|---------|---------|--------|------|--------|---------|----------|-------|
|                 | Solved | Time    | Solved | Time | Solved  | Time    | Solved | Time | Solved | Time    | Solved   | Time  |
| 5               | 100%   | 36      | 100%   | 28   | 100%    | 31      | 100%   | 20   | 100%   | 31      | 100%     | 779   |
| 10              | 100%   | 204     | 100%   | 73   | 100%    | 381     | 100%   | 35   | 100%   | 164     | 100%     | 1,080 |
| 15              | 86%    | 39,701  | 100%   | 148  | 98%     | 67,161  | 100%   | 53   | 100%   | 3,927   | 100%     | 1,450 |
| 20              | 0%     | –       | 100%   | 188  | 0%      | –       | 100%   | 73   | 74%*   | 242,807 | 100%     | 1,777 |
| 25              | 0%     | –       | 100%   | 295  | 0%      | –       | 100%   | 119  | 0%     | –       | 100%     | 1,608 |
| $p_1$           | DPOP   |         | H-DPOP |      | PH-DPOP |         | AFB    |      | ODPOP  |         | ASP-DPOP |       |
|                 | Solved | Time    | Solved | Time | Solved  | Time    | Solved | Time | Solved | Time    | Solved   | Time  |
| 0.4             | 100%   | 1,856   | 100%   | 119  | 100%    | 2,117   | 100%   | 46   | 100%   | 1,819   | 100%     | 1,984 |
| 0.5             | 100%   | 13,519  | 100%   | 120  | 100%    | 19,321  | 100%   | 50   | 100%   | 2,680   | 100%     | 1,409 |
| 0.6             | 94%    | 42,010  | 100%   | 144  | 100%    | 54,214  | 100%   | 51   | 100%   | 3,378   | 100%     | 1,308 |
| 0.7             | 56%    | 66,311  | 100%   | 165  | 88%     | 131,535 | 100%   | 54   | 100%   | 8,063   | 100%     | 1,096 |
| 0.8             | 20%    | 137,025 | 100%   | 164  | 62%     | 247,335 | 100%   | 60   | 100%   | 36,748  | 100%     | 1,073 |

Table 2. Continued

| $ D_i $ | DPOP   |        | H-DPOP |      | PH-DPOP |         | AFB    |      | ODPOP             |         | ASP-DPOP |         |
|---------|--------|--------|--------|------|---------|---------|--------|------|-------------------|---------|----------|---------|
|         | Solved | Time   | Solved | Time | Solved  | Time    | Solved | Time | Solved            | Time    | Solved   | Time    |
| 4       | 100%   | 782    | 100%   | 87   | 100%    | 1,512   | 100%   | 46   | 100%              | 285     | 100%     | 1,037   |
| 6       | 90%    | 28,363 | 100%   | 142  | 98%     | 42,275  | 100%   | 50   | 100%              | 4,173   | 100%     | 1,283   |
| 8       | 14%    | 37,357 | 100%   | 194  | 52%     | 262,590 | 100%   | 60   | 98%               | 71,512  | 100%     | 8,769   |
| 10      | 0%     | –      | 100%   | 320  | 8%      | 354,340 | 100%   | 70   | 78% <sup>†</sup>  | 227,641 | 100%     | 29,598  |
| 12      | 0%     | –      | 100%   | 486  | 0%      | –       | 100%   | 82   | 30% <sup>‡</sup>  | 343,756 | 100%     | 60,190  |
| $p_2$   | DPOP   |        | H-DPOP |      | PH-DPOP |         | AFB    |      | ODPOP             |         | ASP-DPOP |         |
|         | Solved | Time   | Solved | Time | Solved  | Time    | Solved | Time | Solved            | Time    | Solved   | Time    |
| 0.3     | 90%    | 38,114 | 100%   | 464  | 76%     | 189,431 | 100%   | 103  | 84% <sup>§</sup>  | 221,515 | 18%      | 120,114 |
| 0.4     | 86%    | 48,632 | 100%   | 265  | 84%     | 107,986 | 100%   | 71   | 94% <sup>  </sup> | 109,961 | 86%      | 50,268  |
| 0.5     | 94%    | 38,043 | 100%   | 161  | 96%     | 71,181  | 100%   | 57   | 100%              | 14,790  | 92%      | 4,722   |
| 0.6     | 90%    | 31,513 | 100%   | 144  | 98%     | 68,307  | 100%   | 52   | 100%              | 13,519  | 100%     | 1,410   |
| 0.7     | 90%    | 39,352 | 100%   | 128  | 100%    | 49,377  | 100%   | 48   | 100%              | 1,730   | 100%     | 1,059   |
| 0.8     | 92%    | 40,526 | 100%   | 112  | 100%    | 62,651  | 100%   | 57   | 100%              | 1,137   | 100%     | 1,026   |

\*ODPOP cannot solve 13 instances (out of 50 instances) in this experiment in which there are 12 instances unsolved due to timeout and 1 instance unsolved due to memory limitation. <sup>†</sup>ODPOP cannot solve 11 instances (out of 50 instances) in this experiment in which there are 10 instances unsolved due to timeout and 1 instance unsolved due to memory limitation. <sup>‡</sup>ODPOP cannot solve 35 instances (out of 50 instances) in this experiment in which there are 29 instances unsolved due to timeout and 6 instance unsolved due to memory limitation. <sup>§</sup>ODPOP cannot solve 8 instances (out of 50 instances) in this experiment due to timeout. <sup>||</sup>ODPOP cannot solve 3 instances (out of 50 instances) in this experiment due to timeout.

Table 3. Additional experimental results on random graphs

| $ \mathcal{X} $ | AFB    |         | ASP-DPOP |         |
|-----------------|--------|---------|----------|---------|
|                 | Solved | Time    | Solved   | Time    |
| 150             | 100%   | 31,156  | 100%     | 37,862  |
| 200             | 100%   | 117,913 | 100%     | 115,966 |
| 250             | 0%*    | –       | 100%     | 298,361 |

\*AFB cannot solve any instance (out of 50 instances) in this experiment due to timeout.

constraints, unlike ASP-DPOP. On one hand, ODPOP intuitively sends each row of UTIL messages per time on demand and uses optimality criteria to prove optimality without exploring all value assignments for the respective variables. However, these techniques are not as efficient as pruning the search space in ASP-DPOP when the problem becomes more complex. Thus, ODPOP reaches a timeout in most of its unsolvable problems. It is also worth to observe that there are some problems that ODPOP cannot solve due to memory limitations. We attribute this to the fact that ODPOP maintains in its search space *infeasible* value assignments that result in a utility equal to  $-\infty$ , and thus the search space of ODPOP is not as optimized as that of ASP-DPOP.

*Additional experiment results on random graphs:* We claimed earlier that AFB is able to solve more problems and solve them faster than every other algorithm, mainly due to the relative small number of variables in the experiments reported in Table 2. To directly confirm such claim, we extended our experiments on random graphs, by increasing the number of variables (i.e.,  $|\mathcal{X}| \geq 150$ ) and keeping the other parameters to their “default” values (i.e.,  $|\mathcal{A}| = 5, |D_i| = 6, p_1 = 0.6, p_2 = 0.6$ ).<sup>8</sup> The timeout is also set to  $5 \times 10^6$  ms. Table 3 shows the percentage of instances solved (out of 50 instances) and the average simulated runtime (in ms) for the solved instances. The runtime results for DPOP, H-DPOP, PH-DPOP, and ODPOP are not included in Table 3 because they run out of memory in solving all of the problems in this domain.<sup>9</sup> We observe that ASP-DPOP is able to solve more problems than AFB (i.e., when  $|\mathcal{X}| = 250$ ) and solve them faster than AFB when  $|\mathcal{X}| \geq 200$ . We attribute this observation mainly to the large number of variables in this experiment. We also notice that AFB can scale up to solve problems of up to  $|\mathcal{X}| = 200$  (such scalability will not be seen in the experiment on power network problems described below). The main reason is that the problems in our experiment on random graphs are “purely hard” with the default values  $p_1 = 0.6$  and  $p_2 = 0.6$ . This means that the size of the set of complete feasible value assignments, which are complete value assignments that do not result in a utility of  $+\infty$ , is small (about less than 5 in all of the problems in this domain). AFB backtracks much earlier before it can achieves

<sup>8</sup> We thank one of the reviewers for his/her suggestion to have this additional experiment on random graphs.

<sup>9</sup> It is worth to note that H-DPOP runs out of memory while constructing its CDDs in solving all of the problems in this domain.

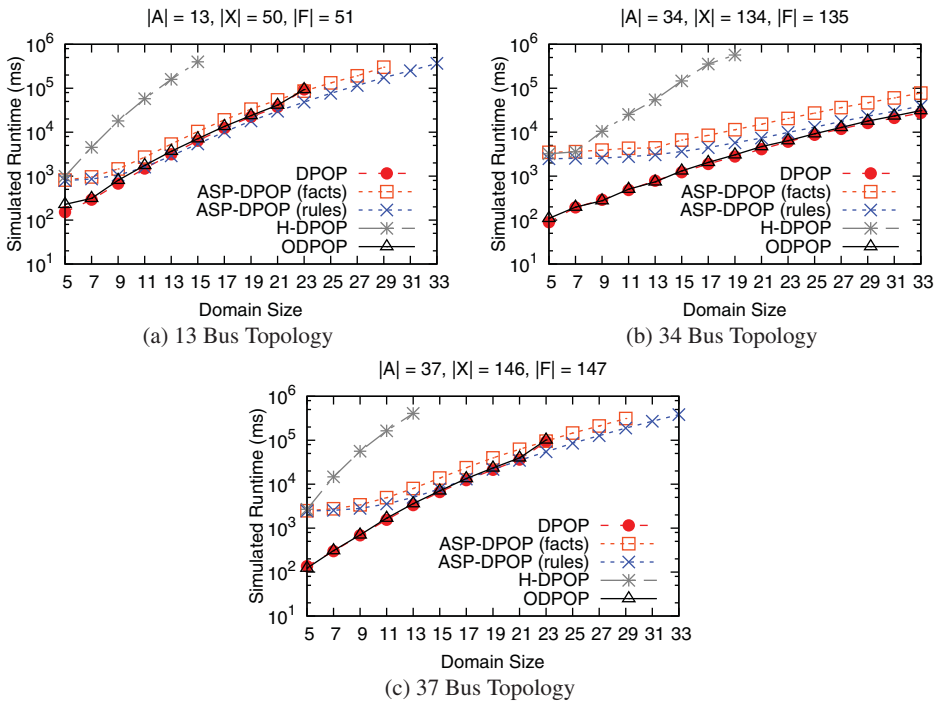


Fig. 5. Runtime results on power network problems. (a) 13 bus topology. (b) 34 bus topology. (c) 37 bus topology.

a complete feasible value assignment. As a result, AFB can solve problems with the number of variables up to 200 before it exceeds the time out threshold.

*Power network problems:* A customer-driven microgrid (CDMG), one possible instantiation of the smart grid problem, has recently been shown to subsume several classical power system sub-problems (e.g., load shedding, demand response, restoration) (Jain *et al.* 2012). In this domain, each agent represents a node with consumption, generation, and transmission preferences, and a global cost function. Constraints include the power balance and no power loss principles, the generation and consumption limits, and the capacity of the power line between nodes. The objective is to minimize a global cost function. CDMG optimization problems are well suited to be modeled with DCOPs due to their distributed nature. Moreover, as some of the constraints in CDMGs (e.g., the power balance principle) can be described in the functional form, they can be exploited by ASP-DPOP (rules). For this reason, both “ASP-DPOP (facts)” and “ASP-DPOP (rules)” are used in this domain.

We use three network topologies defined using the IEEE standards (IEEE Distribution Test Feeders 2014) and vary the domain size of the generation, load, and transmission variables of each agent from 5 to 31. The timeout is set to  $10^6$  ms. Figure 5 summarizes the runtime results. As the utilities are generated following predefined rules (Gupta *et al.* 2013), we also show the results for ASP-DPOP (rules). Furthermore, we omit results for PH-DPOP because they have identical runtime—

Table 4. Message size results on power network problems

| $ D_i $                       | 13 bus topology |         |         |         | 34 bus topology |        |         |         |
|-------------------------------|-----------------|---------|---------|---------|-----------------|--------|---------|---------|
|                               | 5               | 7       | 9       | 11      | 5               | 7      | 9       | 11      |
| H-DPOP                        | 6,742           | 30,604  | 97,284  | 248,270 | 1,437           | 4,685  | 11,617  | 24,303  |
| DPOP                          | 3,125           | 16,807  | 59,049  | 161,051 | 625             | 2,401  | 6,561   | 14,641  |
| ODPOP                         | 6               | 6       | 6       | 6       | 5               | 5      | 5       | 5       |
| ASP-DPOP                      | 10              | 14      | 18      | 22      | 10              | 14     | 18      | 22      |
| 37 bus topology               |                 |         |         |         |                 |        |         |         |
| $ D_i $                       | 5               | 7       | 9       | 11@     |                 |        |         |         |
| H-DPOP                        | 6,742           | 30,604  | 97,284  | 248,270 |                 |        |         |         |
| DPOP                          | 3,125           | 16,807  | 59,049  | 161,051 |                 |        |         |         |
| ODPOP                         | 6               | 6       | 6       | 6       |                 |        |         |         |
| ASP-DPOP                      | 10              | 14      | 18      | 22      |                 |        |         |         |
| (a) Largest UTIL message size |                 |         |         |         |                 |        |         |         |
| $ D_i $                       | 13 bus topology |         |         |         | 34 bus topology |        |         |         |
|                               | 5               | 7       | 9       | 11      | 5               | 7      | 9       | 11      |
| H-DPOP                        | 19,936          | 79,322  | 236,186 | 579,790 | 20,810          | 57,554 | 130,050 | 256,330 |
| DPOP                          | 9,325           | 43,687  | 143,433 | 375,859 | 9,185           | 29,575 | 73,341  | 153,923 |
| ODPOP                         | 391             | 1,430   | 6,281   | 11,979  | 2,197           | 4,122  | 12,124  | 12,870  |
| ASP-DPOP                      | 120             | 168     | 216     | 264     | 330             | 462    | 594     | 726     |
| 37 bus topology               |                 |         |         |         |                 |        |         |         |
| $ D_i $                       | 5               | 7       | 9       | 11      |                 |        |         |         |
| H-DPOP                        | 38,689          | 133,847 | 363,413 | 836,167 |                 |        |         |         |
| DPOP                          | 17,665          | 71,953  | 215,793 | 531,025 |                 |        |         |         |
| ODPOP                         | 1,896           | 5,572   | 18,981  | 28,285  |                 |        |         |         |
| ASP-DPOP                      | 360             | 504     | 648     | 792     |                 |        |         |         |
| (b) Total UTIL message size   |                 |         |         |         |                 |        |         |         |

the amount of information used to prune the search space is identical for both algorithms in this domain. We also measure the size of UTIL messages, where we use the number of values in the message as units, and the intra-agent UTIL messages (i.e., messages are passed between virtual agents that belong to the same real agent) are accounted for fair comparison. Table 4 tabulates the results. We did not measure the size of VALUE messages since they are significantly smaller than UTIL messages. It is also worth to report that the number of UTIL messages that FRODO produces (discounting all intra-agent UTIL messages) is equal to the number of UTIL messages that ASP-DPOP produced in all power network problems in our experiments.

The results in Figure 5 are consistent with those shown earlier (except for AFB)—ASP-DPOP is slower than DPOP and ODPOP when the domain size is small, but it is able to solve more problems than DPOP and ODPOP. We observe that, in

Figure 5(b), DPOP is consistently faster than ASP-DPOP and is able to solve the same number of problems as ASP-DPOP. It is because the highest constraint arity in 34 bus topology is 5, while it is 6 in 13 and 37 bus topologies. Unlike in random graphs, H-DPOP is slower than the other algorithms in these problems. The reason is that the constraint arity in these problems is larger and the expensive operations on CDDs grow exponentially with the arity. We also observe that ASP-DPOP (rules) is faster than ASP-DPOP (facts). The reason is that the former is able to exploit the interdependencies between constraints to prune the search space. Additionally, ASP-DPOP (rules) can solve more problems than ASP-DPOP (facts). The reason is that the former requires less memory since it prunes a larger search space and, thus, ground fewer facts.

The runtime results for AFB are not included in Figure 5, since AFB exceeds the timeout in solving all of the problems in this domain; this contrasts to the results shown earlier for random graphs. The main reason is that the number of variables in the power network problems is large (i.e.,  $|\mathcal{X}|$  are 50, 134, and 146 in 13, 34, and 37 bus topologies, respectively in Fig. 5).

Finally, both versions of ASP-DPOP require smaller messages than both H-DPOP and DPOP. The reason for the former is that the CDD data structure of H-DPOP is significantly more complex than that of ASP-DPOP. The reason for the latter is that ASP-DPOP prunes portions of the search space while DPOP did not. In addition, since ASP-DPOP does not transform DCOP problems with multiple variables per agent to corresponding ones with one variable per agent, ASP-DPOP is able to exploit significantly more the interdependencies between constraints to prune the search space. Moreover, we can see that the largest UTIL message sizes in ODPOP are smaller than those of ASP-DPOP, but the total UTIL message sizes in ODPOP are larger than those of ASP-DPOP. The reason is that ODPOP sends only linear size message, but it needs to send many messages on demand.

### 5.6 Discussions on ASP-DPOP

ASP-DPOP has been shown to be competitive with other algorithms in solving DCOPs in our experimental results. The benefits of using ASP-DPOP are accomplished by having ASP as its foundation. We will illustrate here the two main advantages of making use of ASP within ASP-DPOP:

- (1) The use of the highly expressive ASP language to encode constraints in DCOPs.
- (2) The ability to harness the highly optimized ASP grounder and solver to prune the search space based on hard constraints.

In the rest of this section, we further discuss these advantages and relate them to the observations drawn from the experiments. These considerations are followed by a discussion of how ASP-DPOP alleviates the simplifying assumption of having a single variable per agent. Finally, at the end of this section, we analyze the privacy loss of ASP-DPOP.

The first advantage of using ASP within ASP-DPOP comes from the ability to use a very expressive logic language to encode the constraints in a DCOP. ASP-



DPOP can represent constraint utilities as an implicit function instead of explicitly enumerating them. Thus, ASP-DPOP is particularly suitable to encode DCOPs whose constraint utilities are large and evaluated via implicit functions of the variables in their scopes (e.g., power network problems, smart grid problems). This can be seen clearly via Example 15.

*Example 15*

Let us consider a constraint  $f$  representing the power loss principle in a power network problem, where  $scp(f) = \{x_{1 \rightarrow 2}, x_{2 \rightarrow 1}\}$  in which the domains of the variables  $x_{1 \rightarrow 2}$  and  $x_{2 \rightarrow 1}$  are  $D_{1 \rightarrow 2} = [0, 2]$  and  $D_{2 \rightarrow 1} = [-2, 0]$ , respectively. Intuitively, the variable  $x_{i \rightarrow j}$ , where  $i, j \in \{1, 2\}, i \neq j$ , indicates the amount of power that node  $i$  transfers to (receives from) node  $j$  if  $x_{i \rightarrow j} \geq 0$  (resp.  $x_{i \rightarrow j} < 0$ ). For example,  $x_{1 \rightarrow 2} = 1$  means that the node 1 transfers 1 unit of power to the node 2, and  $x_{2 \rightarrow 1} = -1$  specifies that the node 2 receives 1 unit of power from the node 1. By the power loss principle, if there is no loss, the amount of power transferred from one node is equal to the amount of power received in the other node (i.e.,  $x_{i \rightarrow j} + x_{j \rightarrow i} = 0$ ). However, if there is loss (i.e.,  $x_{i \rightarrow j} + x_{j \rightarrow i} \neq 0$ ), we assume that the cost (utility) of the power transmission is evaluated to be two times greater than the power unit loss. Formally, the utility of the constraint  $f$  is given implicitly as a function:

$$f(x_{1 \rightarrow 2}, x_{2 \rightarrow 1}) = 2 \times |x_{1 \rightarrow 2} + x_{2 \rightarrow 1}|. \tag{5.1}$$

| $x_{1 \rightarrow 2}$ | $x_{2 \rightarrow 1}$ | Utilities |
|-----------------------|-----------------------|-----------|
| 2                     | -2                    | 0         |
| 2                     | -1                    | 2         |
| 2                     | 0                     | 4         |
| 1                     | -2                    | 2         |
| 1                     | -1                    | 0         |
| 1                     | 0                     | 2         |
| 0                     | -2                    | 4         |
| 0                     | -1                    | 2         |
| 0                     | 0                     | 0         |

(a) Explicit representation as a utility table

$$value(x_{1 \rightarrow 2}, 0..2) \leftarrow \tag{5.2}$$

$$value(x_{2 \rightarrow 1}, -2..0) \leftarrow \tag{5.3}$$

$$f(2 * |V_1 + V_2|, V_1, V_2) \leftarrow value(x_{1 \rightarrow 2}, V_1), value(x_{2 \rightarrow 1}, V_2). \tag{5.4}$$

(b) Implicit representation as an answer set program

Fig. 6. Different encodings of constraint  $f$  in Example 15.

Figure 6(a) enumerates all the utilities of the constraint  $f$  explicitly in a utility table, and Figure 6(b) presents an answer set program that models implicitly those utilities. We can see that while the utility table has nine rows (i.e., the domain sizes of  $x_{1 \rightarrow 2}$  and  $x_{2 \rightarrow 1}$  are 3), the answer set program consists of only two facts and one rule. If the domain sizes of  $x_{1 \rightarrow 2}$  and  $x_{2 \rightarrow 1}$  are 1,000 (e.g.,  $D_{1 \rightarrow 2} = [0, 999]$  and  $D_{2 \rightarrow 1} = [-999, 0]$ ), the utility table would have **1,000<sup>2</sup>** rows whereas the answer set program modeling implicitly such the same utilities still has **two facts** and **one rule** that are similar to ones in Figure 6(b)—i.e., it only updates the two facts (5.2)

and (5.3) as follows:

$$\text{value}(x_{1 \rightarrow 2}, 0..999) \leftarrow \quad (5.5)$$

$$\text{value}(x_{2 \rightarrow 1}, -999..0) \leftarrow \quad (5.6)$$

As a consequence, using ASP within ASP-DPOP to encode DCOPs makes programs much more concise and compact. The encoding is declarative and can be easily extended and modified. Moreover, such encoding does not depend on the implementation of the algorithms (e.g., DPOP or H-DPOP), making programs more flexible and understandable. Specifically, if we change the algorithm to solve a DCOP, the CM needs to be changed following the new algorithm, yet the SM remains the same. In contrast, using imperative programming techniques, the “ad-hoc” implementation that is employed within each local solver might require different encodings of DCOPs for different used algorithms and different propagators for different types of constraints. For example, H-DPOP implementation needs a different data structure from DPOP implementation to deal with hard constraints.

The second advantage of using ASP as the foundation of ASP-DPOP is to harness the highly optimized ASP grounders and solvers to prune the search space, especially in the handling of hard constraints. As an example, consider the power network problem whose objective is to minimize a global cost function.<sup>10</sup> A DCOP that encodes such type of power network problems can be formulated in terms of cost-as-utility *minimization* rather than reward-as-utility *maximization*. Thus, in this formulation, the value assignments resulting in an infinite utility (i.e.,  $+\infty$ ) should not be included in any DCOP solution; such value assignments are redundant and should be pruned. Example 16 shows how effectively an ASP grounder can prune the search space.

#### Example 16

Consider a simple power network problem, where the aggregated cost needs to be minimized. The problem has two nodes (nodes 1 and 2). Let us assume that agent  $a_1$  and agent  $a_2$ , which are the node 1 and the node 2, own the variables  $x_{1 \rightarrow 2}$  and  $x_{2 \rightarrow 1}$ , respectively. These are described in Example 15. The problem has one constraint  $f$  representing the power loss principle, analogously to what described in Example 15. The only difference is that we do not allow losses in power transfers (i.e., if there is a loss, the corresponding cost is  $+\infty$ ). Thus, the utility (cost) of the constraint  $f$  now is evaluated as

$$f(x_{1 \rightarrow 2}, x_{2 \rightarrow 1}) = \begin{cases} 2 \times x_{1 \rightarrow 2} & \text{if } x_{1 \rightarrow 2} + x_{2 \rightarrow 1} = 0 \\ +\infty & \text{otherwise} \end{cases} \quad (5.7)$$

Figure 7 presents the ASP program<sup>11</sup> to compute the UTIL message sent from the agent 2 to the agent 1, assuming that the agent 1 is the root of the respective pseudo-tree (i.e., the separator set of the agent 2 is  $sep_2 = \{x_{1 \rightarrow 2}\}$ ). It is important to observe

<sup>10</sup> The previous formalization of ASP-DPOP focuses on maximizing the cost function; the switch to minimization problems requires trivial changes to the design of ASP-DPOP.

<sup>11</sup> #sup is a special constant representing the largest possible value in the ASP language.

that, since the objective is to minimize a global cost function, the ASP in Figure 7 is produced differently from the one that is generated by `generate_UTILASP/2` described in 3.4.4. Specifically, the differences are the following:

- The predicates of the form `table_min_` $a_i$  are used instead of ones of the form `table_max_` $a_i$ ;
- $U = \#min\{\dots\}$  rather than  $U = \#max\{\dots\}$  in (3.35) (i.e., computing the minimal utilities for each value combination of variables in the separator list); and
- The conditions  $V_{r_1} \neq \#sup, \dots, V_{r_k} \neq \#sup$  are used instead of  $V_{r_1} \neq \#inf, \dots, V_{r_k} \neq \#inf$  in (3.34) (i.e., atoms of the form `table_row_` $a_i(u, v_{s_1}, \dots, v_{s_k})$  where  $u = \#sup$  (i.e., the respective utilities are  $+\infty$ ) are not produced).

As a consequence, the encoded UTIL messages consist of facts of the forms `table_min_` $a_i$  (instead of `table_max_` $a_i$ ) and `table_info`.

$$f(4, 2, -2) \leftarrow \quad (5.8)$$

$$f(2, 1, -1) \leftarrow \quad (5.9)$$

$$f(0, 0, 0) \leftarrow \quad (5.10)$$

$$f(\#sup, 0, -2) \leftarrow \quad (5.11)$$

$$f(\#sup, 1, -2) \leftarrow \quad (5.12)$$

$$f(\#sup, 0, -1) \leftarrow \quad (5.13)$$

$$f(\#sup, 2, -1) \leftarrow \quad (5.14)$$

$$f(\#sup, 1, 0) \leftarrow \quad (5.15)$$

$$f(\#sup, 2, 0) \leftarrow \quad (5.16)$$

$$value(x_{1 \rightarrow 2}, 0..2) \leftarrow \quad (5.17)$$

$$value(x_{2 \rightarrow 1}, -2..0) \leftarrow \quad (5.18)$$

$$\begin{aligned} table\_row\_a_2(U, X_{1 \rightarrow 2}) \leftarrow & f(V_0, X_{1 \rightarrow 2}, X_{2 \rightarrow 1}), \\ & V_0 \neq \#sup, \\ & U = V_0. \end{aligned} \quad (5.19)$$

$$\begin{aligned} table\_min\_a_2(U, X_{1 \rightarrow 2}) \leftarrow & value(x_{1 \rightarrow 2}, X_{1 \rightarrow 2}), \\ & table\_row\_a_2(\_, X_{1 \rightarrow 2}), \\ & U = \#min\{V : table\_row\_a_2(V, X_{1 \rightarrow 2})\}. \end{aligned} \quad (5.20)$$

Fig. 7. ASP to compute UTIL message in example 16.

The nine facts (5.8)–(5.16) enumerate all utilities of the constraint  $f$  in which the six facts (5.11)–(5.16) are redundant since their corresponding utilities are  $+\infty$ . With DPOP, the total size of the search space for computing its UTIL message is 9, which corresponds to the nine facts (5.8)–(5.16), since DPOP does not do pruning. However, with ASP-DPOP, the corresponding total size of the search space is 3 since GRINGO, due to the condition  $V_0 \neq \#sup$ , grounds the rule (5.19) into only three

facts:

$$\text{table\_row\_}a_2(4,2) \leftarrow \quad (5.21)$$

$$\text{table\_row\_}a_2(2,1) \leftarrow \quad (5.22)$$

$$\text{table\_row\_}a_2(0,0) \leftarrow \quad (5.23)$$

and an ASP solver will use these facts to generate the predicates `table_min_a2` based on the rule (5.20). The difference between the sizes of the search spaces of ASP-DPOP and DPOP are greater as the domain sizes of variables increase. For example, if the domain sizes of  $x_{1 \rightarrow 2}$  and  $x_{2 \rightarrow 1}$  are 1,000, the total search space of DPOP is  $1,000^2$ , while the total search space of ASP-DPOP is just **1,000**.

As a consequence, and as clear from our experiments, ASP-DPOP is able to prune a significant portion of the search space, thanks to hard constraints, whereas DPOP does not. Moreover, as seen in Example 16, the size of the search space pruned increases as the complexity of the instance grows (i.e., increasing  $|\mathcal{X}|$ ,  $|D_i|$ ,  $p_1$ , or  $p_2$ ). Thus, ASP-DPOP is able to solve more problems than DPOP and is faster than DPOP when the problem becomes more complex.

The pruning power of the ASP grounders and solvers enables also the generation of smaller UTIL messages in ASP-DPOP than those generated by DPOP. Let us consider a UTIL message  $M$  sent from an agent  $a_i$  to an agent  $a_j$ . A value assignment of variables in  $sep_i$  is *admissible* if its corresponding optimal sum of utilities in the subtree rooted at  $a_i$  is different than  $-\infty$ .<sup>12</sup> In DPOP,  $M$  consists of a utility, which is optimal, for each value assignment of variables in  $sep_i$  (including both admissible and inadmissible value assignments). However,  $M$  in ASP-DPOP consists of a utility, which is optimal and different from  $-\infty$ , for only each admissible value assignment of variables in  $sep_i$ . This is because such inadmissible value assignments will not be included in any DCOP solution (i.e., otherwise the global cost is  $-\infty$ ).

We will not discuss in-depth technically what algorithms and computations are implemented within modern ASP grounders to optimize the grounding process, since they are beyond the scope of this paper. Readers who are interested in such algorithms and computations can find further information in Gebser *et al.* (2012) and Kaufmann *et al.* (2016). It is important to notice that such computations (e.g., for removing unnecessary rules and for omitting rules whose bodies cannot be satisfied) consume memory, take time, and are not trivial. Therefore, for DCOP problems with low constraint tightness, the runtime and memory that are used for those computations will dominate the runtime and memory that are saved from pruning the search space (e.g., see the row  $p_2 = 0.3$  in Table 2). This also explains why ASP-DPOP is slower than DPOP when the problem becomes less complex (i.e., decreasing  $|\mathcal{X}|$ ,  $|D_i|$ ,  $p_1$ , or  $p_2$ ). Specifically, from the trend while decreasing  $p_2$  in Table 2, ASP-DPOP will not be able to compete with DPOP for cases where  $p_2 \leq 0.3$ .

The fact that ASP-DPOP solves DCOP problems with multiple variables per agent directly, without transforming them to problems with one variable per agent, deserves

<sup>12</sup> Or  $+\infty$  for minimization problems.

some discussions. It is easy to see that ASP-DPOP agents need to consider more variables and thus more constraints. As a result, there are more interdependencies between constraints for ASP-DPOP to exploit. If the constraint tightness is high, the size of the search space pruned increases significantly. This can be seen in our power network experiment. On the other hand, dealing with more variables and more constraints also increases the search space. Therefore, if the constraint tightness does not provide sufficient pruning, the portion of the search space pruned does not properly balance the increase in the size of the search space; this may lead ASP-DPOP to require more memory than DPOP in solving such problems. This situation can be seen in experimental results on random graphs (i.e., decreasing  $p_2$ ). Solving DCOPs with multiple variables per agent without transforming them to problems with a single variable per agent was also investigated in Fioretto *et al.* (2016).

Maintaining privacy is a fundamental motivation for the use of DCOP. A detailed analysis of privacy loss in DCOP for some existing DCOP algorithms, including DPOP, can be found in Greenstadt *et al.* (2006). For ASP-DPOP, it is not difficult to realize that DPOP and ASP-DPOP have the same privacy loss. The reason is that the content of UTIL messages (resp. VALUE messages) in DPOP—that are given under the tabular form (which are similar to those given under multi-dimensional matrix form)—is identical to the content of the UTIL messages (resp. VALUE messages) in ASP-DPOP—that are given in facts form. In fact, anything that is inferred from the fact form (in UTIL and VALUE messages of ASP-DPOP) can be inferred from the tabular form (in the respective messages of DPOP), and vice versa anything is inferred from tabular form can be inferred from fact form as well.

## 6 Related work

The use of declarative programs, specifically logic programs, for reasoning in multi-agent domains is not new. Starting with some seminal papers (Kowalski and Sadri 1999), various authors have explored the use of several different flavors of logic programming, such as normal logic programs and abductive logic programs, to address cooperation between agents (Sadri and Toni 2003; Kakas *et al.* 2004; De Vos *et al.* 2005; Gelfond and Watson 2007). Some proposals have also explored the combination between constraint programming, logic programming, and formalization of multi-agent domains (Vlahavas 2002; Dovier *et al.* 2010a, 2010b, 2013). Logic programming has been used in modeling multi-agent scenarios involving agents knowledge about other's knowledge (Baral *et al.* 2010), computing models in the logics of knowledge (Pontelli *et al.* 2010), multi-agent planning (Son *et al.* 2009) and formalizing negotiation (Sakama *et al.* 2011). ASP-DPOP is similar to the last two applications in that (i) it can be viewed as a collection of agent programs; (ii) it computes solutions using an ASP solver; and (iii) it uses message passing for agent communication. A key difference is that ASP-DPOP solves multi-agent problems formulated as constraint-based models, while the other applications solve problems formulated as decision-theoretic and game-theoretic models.

Researchers have also developed a framework that integrates declarative techniques with off-the-shelf constraint solvers to partition large constraint optimization problems into smaller subproblems and solve them in parallel (Liu *et al.* 2012). In contrast, DCOPs are problems that are naturally distributed and cannot be arbitrarily partitioned.

ASP-DPOP is able to exploit problem structure by propagating hard constraints and using them to prune the search space efficiently. This reduces the memory requirement of the algorithm and improves the scalability of the system. Existing DCOP algorithms that also propagate hard and soft constraints to prune the search space include H-DPOP that propagates exclusively hard constraints (Kumar *et al.* 2008), BrC-DPOP that propagates branch consistency (Fioretto *et al.* 2014), and variants of BnB-ADOPT (Yeoh *et al.* 2010; Gutierrez *et al.* 2011; Gutierrez and Meseguer 2012b) that maintains soft-arc consistency (Bessiere *et al.* 2012; Gutierrez and Meseguer 2012a; Gutierrez *et al.* 2013). A key difference is that these algorithms require algorithm developers to explicitly implement the ability to reason about the hard and soft constraints and propagate them efficiently. In contrast, ASP-DPOP capitalizes on general purpose ASP solvers to do so.

## 7 Conclusions and Future work

In this paper, we explored the benefits of using logic programming techniques as a platform to provide complete solutions of DCOPs. Our proposed logic programming-based algorithm, ASP-DPOP, is able to solve more problems and solve them faster than DPOP, its imperative programming counterpart. Aside from the ease of modeling, each agent in ASP-DPOP also capitalizes on highly efficient ASP solvers to automatically exploit problem structure (e.g., prune the search space using hard constraints). Experimental results show that ASP-DPOP is faster and can scale to larger problems than a version of H-DPOP (i.e., PH-DPOP) that maintains the level of privacy similar to that of ASP-DPOP. These results highlight the strengths of a declarative programming paradigm, where explicit model-specific pruning rules are not necessary. In conclusion, we believe that this work contributes to the DCOP community, where we show that the declarative programming paradigm is a promising new direction of research for DCOP researchers, as well as the ASP community, where we demonstrate the applicability of ASP to solve a wide array of multi-agent problems that can be modeled as DCOPs.

In future work, we will explore two directions to deepen the use of logic programming in solving DCOPs:

- *Logic programming under different semantics*: We will consider the advantages of other logic programming paradigms in solving DCOPs. One possibility is to use Constraint Logic Programming (CLP) (Jaffar and Maher 1994) instead of ASP. Since CLP is a merger of two declarative paradigms—constraint solving and logic programming—it seems well-suited to solve DCOPs. A preliminary investigation (Le *et al.* 2014) has shown that this technique can dramatically decrease run time.

- *Different representation of messages*: We observe that the messages used in DPOP, and even ASP-DPOP, are represented explicitly—i.e., they are multi-dimensional matrices in DPOP and facts in ASP-DPOP. One of the reasons for this is that each agent performs the inference process for its subtree, enumerates explicitly all the results, and sends them to other agents. We are interested in investigating algorithms where agents coordinate with others via messages that are logic programs (e.g., ASP or CLP clauses). Specifically, in such an algorithm, each agent does the inference partially, for some specific interesting value assignment, and without enumerating all results. The rest of the computation will be encoded as logic programs and passed to other agents. Some agent who performs the complete inference process will propagate the search space based on the rules in the received messages as logic programs. We believe this will reduce the search space and the run time.

### Acknowledgement

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the sponsoring organizations, agencies, or the U.S. government. We would like to thank Akshat Kumar for sharing with us his implementation of H-DPOP.

### References

- BARAL, C. 2003. *Knowledge Representation, Reasoning, and Declarative Problem Solving with Answer Sets*. Cambridge University Press, Cambridge, MA.
- BARAL, C., GELFOND, G., PONTELLI, E. AND SON, T. C. 2010. Modeling multi-agent scenarios involving agents knowledge about other's knowledge using ASP. In *Proc. of AAMAS*, 259–266.
- BESSIERE, C., GUTIERREZ, P. AND MESEGUER, P. 2012. Including soft global constraints in DCOPs. In *Proc. of CP*, 175–190.
- CARLSSON, M. et al. 2015. *SICStus Prolog User's Manual*. Swedish Institute of Computer Science.
- CITRIGNO, S., EITER, T., FABER, W., GOTTLÖB, G., KOCH, C., LEONE, N., MATEIS, C., PFEIFER, G. AND SCARCELLO, F. 1997. The dlvs system: Model generator and application frontends. In *Proc. of Workshop on Logic Programming*, 128–137.
- DE VOS, M., CRICK, T., PADGET, J. A., BRAIN, M., CLIFFE, O. AND NEEDHAM, J. 2005. LAIMA: A multi-agent platform using ordered choice logic programming. In *Proc. of DALI*.
- DECHTER, R. 2003. *Constraint Processing*. Elsevier/Morgan Kaufmann.
- DOVIER, A., FORMISANO, A. AND PONTELLI, E. 2010a. An investigation of multi-agent planning in CLP. *Fundamentae Informatica* 105, 1–2, 79–103.
- DOVIER, A., FORMISANO, A. AND PONTELLI, E. 2010b. Multivalued action languages with constraints in CLP(FD). *Theory and Practice of Logic Programming* 10, 2, 167–235.
- DOVIER, A., FORMISANO, A. AND PONTELLI, E. 2013. Autonomous agents coordination: Action languages meet CLP() and Linda. *Theory and Practice of Logic Programming* 13, 2, 149–173.
- ERDÖS, P. AND RÉNYI, A. 1959. On random graphs I. *Publicationes Mathematicae Debrecen* 6, 290.

- FARINELLI, A., ROGERS, A., PETCU, A. AND JENNINGS, N. 2008. Decentralised coordination of low-power embedded devices using the Max-Sum algorithm. In *Proc. of AAMAS*, 639–646.
- FIORETTO, F., CAMPEOTTO, F., DA RIN FIORETTO, L., YEOH, W. AND PONTELLI, E. 2014. GD-Gibbs: A GPU-based sampling algorithm for solving distributed constraint optimization problems (Extended Abstract). In *Proc. of AAMAS*.
- FIORETTO, F., LE, T., YEOH, W., PONTELLI, E. AND SON, T. C. 2014. Improving DPOP with branch consistency for solving distributed constraint optimization problems. In *Proc. of CP*.
- FIORETTO, F., YEOH, W. AND PONTELLI, E. 2016. Multi-variable agents decomposition for dcops. In *Proc. of 30th AAAI Conference on Artificial Intelligence*, Phoenix, Arizona, USA, February 12–17, 2016, 2480–2486.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B. AND SCHAUB, T. 2012. *Answer Set Solving in Practice*. Morgan and Claypool Publishers.
- GEBSER, M., KAUFMANN, B., KAMINSKI, R., OSTROWSKI, M., SCHAUB, T. AND SCHNEIDER, M. 2011. Potassco: The potsdam answer set solving collection. *AI Communications* 24, 2 (Apr.), 107–124.
- GEBSER, M., KAUFMANN, B., NEUMANN, A. AND SCHAUB, T. 2007. Clasp: A conflict-driven answer set solver. In *Proc. of LPNMR*, 260–265.
- GELFOND, G. AND WATSON, R. 2007. Modeling cooperative multi-agent systems. In *Proc. of ASP Workshop*.
- GELFOND, M. AND KAHL, Y. 2014. *Knowledge Representation, Reasoning, and the Design of Intelligent Agents*. Cambridge University Press.
- GELFOND, M. AND LIFSCHITZ, V. 1990. Logic programs with classical negation. In *Proc. of ICLP*, 579–597.
- GERSHMAN, A., MEISELS, A. AND ZIVAN, R. 2009. Asynchronous forward-bounding for distributed COPs. *Journal of Artificial Intelligence Research* 34, 61–88.
- GREENSTADT, R., PEARCE, J. P. AND TAMBE, M. 2006. Analysis of privacy loss in distributed constraint optimization. In *Proc. of 21st National Conference on Artificial Intelligence and the 18th Innovative Applications of Artificial Intelligence Conference*. July 16–20, 2006, Boston, MA, USA, 647–653.
- GUPTA, S., JAIN, P., YEOH, W., RANADE, S. AND PONTELLI, E. 2013. Solving customer-driven microgrid optimization problems as DCOPs. In *Proc. of Distributed Constraint Reasoning Workshop*, 45–59.
- GUTIERREZ, P., LEE, J., LEI, K. M., MAK, T. AND MESEGUER, P. 2013. Maintaining soft arc consistencies in BnB-ADOPT<sup>+</sup> during search. In *Proc. of CP*, 365–380.
- GUTIERREZ, P. AND MESEGUER, P. 2012a. Improving BnB-ADOPT<sup>+</sup>-AC. In *Proc. of AAMAS*, 273–280.
- GUTIERREZ, P. AND MESEGUER, P. 2012b. Removing redundant messages in n-ary BnB-ADOPT. *Journal of Artificial Intelligence Research* 45, 287–304.
- GUTIERREZ, P., MESEGUER, P. AND YEOH, W. 2011. Generalizing ADOPT and BnB-ADOPT. In *Proc. of IJCAI*, 554–559.
- HAMADI, Y., BESSIÈRE, C. AND QUINQUETON, J. 1998. Distributed intelligent backtracking. In *Proc. of ECAI*, 219–223.
- IEEE DISTRIBUTION TEST FEEDERS. 2014. URL: <http://ewh.ieee.org/soc/pes/dsacom/testfeeders/> [Accessed on: 29/07/2014].
- JAFFAR, J. AND MAHER, M. J. 1994. Special issue: Ten years of logic programming constraint logic programming: A survey. *The Journal of Logic Programming* 19, 503–581.
- JAIN, P., GUPTA, S., RANADE, S. AND PONTELLI, E. 2012. Optimum operation of a customer-driven microgrid: A comprehensive approach. In *Proc. of PEDES*.



- KAKAS, A., TORRONI, P. AND DEMETRIOU, N. 2004. Agent Planning, negotiation and control of operation. In *Proc. of ECAI*.
- KAUFMANN, B., LEONE, N., PERRI, S. AND SCHAUB, T. 2016. Grounding and solving in answer set programming. *AI Magazine* 37, 3, 25–32.
- KOWALSKI, R. AND SADRI, F. 1999. Logic programming towards multi-agent systems. *Annals of Mathematics and Artificial Intelligence* 25, 3–4, 391–419.
- KUMAR, A., FALTINGS, B. AND PETCU, A. 2009. Distributed constraint optimization with structured resource constraints. In *Proc. of AAMAS*, 923–930.
- KUMAR, A., PETCU, A. AND FALTINGS, B. 2008. H-DPOP: Using hard constraints for search space pruning in DCOP. In *Proc. of AAAI*, 325–330.
- LASS, R., KOPENA, J., SULTANIK, E., NGUYEN, D., DUGAN, C., MODI, P. AND REGLI, W. 2008. Coordination of first responders under communication and resource constraints (Short Paper). In *Proc. of AAMAS*, 1409–1413.
- LE, T., PONTELLI, E., SON, T. C. AND YEOH, W. 2014. Logic and constraint logic programming for distributed constraint optimization. *Technical Communications of the Thirtieth International Conference on Logic Programming (ICLP' 14)*. Theory and Practice of Logic Programming, Online Supplement.
- LE, T., SON, T. C., PONTELLI, E. AND YEOH, W. 2015. Solving distributed constraint optimization problems with logic programming. In *Proc. of AAAI*.
- LÉAUTÉ, T. AND FALTINGS, B. 2011. Coordinating logistics operations with privacy guarantees. In *Proc. of IJCAI*, 2482–2487.
- LÉAUTÉ, T., OTTENS, B. AND SZYMANEK, R. 2009. FRODO 2.0: An open-source framework for distributed constraint optimization. In *Proc. of Distributed Constraint Reasoning Workshop*, 160–164.
- LIU, C., REN, L., LOO, B. T., MAO, Y. AND BASU, P. 2012. Cologne: A declarative distributed constraint optimization platform. *Proc. of VLDB Endowment* 5, 8, 752–763.
- MAHESWARAN, R., PEARCE, J. AND TAMBE, M. 2004. Distributed algorithms for DCOP: A graphical game-based approach. In *Proc. of PDCS*, 432–439.
- MAHESWARAN, R., TAMBE, M., BOWRING, E., PEARCE, J. AND VARAKANTHAM, P. 2004. Taking DCOP to the real world: Efficient complete solutions for distributed event scheduling. In *Proc. of AAMAS*, 310–317.
- MAILLER, R. AND LESSER, V. 2004. Solving distributed constraint optimization problems using cooperative mediation. In *Proc. of AAMAS*, 438–445.
- MAREK, V. AND TRUSZCZYŃSKI, M. 1999. Stable models and an alternative logic programming paradigm. In *The Logic Programming Paradigm: A 25-year Perspective*, Apt, K., Marek, V.W., Truszczyński, M., Warren, D.S., Eds. Springer, Berlin, 375–398.
- MODI, P., SHEN, W.-M., TAMBE, M. AND YOKOO, M. 2005. ADOPT: Asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence* 161, 1–2, 149–180.
- NGUYEN, D. T., YEOH, W. AND LAU, H. C. 2013. Distributed Gibbs: A memory-bounded sampling-based DCOP algorithm. In *Proc. of AAMAS*, 167–174.
- NIEMELÄ, I. 1999. Logic programming with stable model semantics as a constraint programming paradigm. *Annals of Mathematics and Artificial Intelligence* 25, 3–4, 241–273.
- OTTENS, B., DIMITRAKAKIS, C. AND FALTINGS, B. 2012. DUCT: An upper confidence bound approach to distributed constraint optimization problems. In *Proc. of AAAI*, 528–534.
- PETCU, A. 2009. *A Class of Algorithms for Distributed Constraint Optimization*. Frontiers in Artificial Intelligence and Applications, vol. 194. IOS Press.
- PETCU, A. AND FALTINGS, B. 2005a. A scalable method for multiagent constraint optimization. In *Proc. of IJCAI*, 1413–1420.

- PETCU, A. AND FALTINGS, B. 2005b. Superstabilizing, fault-containing multiagent combinatorial optimization. In *Proc. of AAAI*, 449–454.
- PETCU, A. AND FALTINGS, B. 2006. ODPOP: An algorithm for open/distributed constraint optimization. In *Proc. of 21st National Conference on Artificial Intelligence and the 18th Innovative Applications of Artificial Intelligence Conference*. July 16–20, 2006, Boston, Massachusetts, USA, 703–708.
- PETCU, A. AND FALTINGS, B. 2007. MB-DPOP: A new memory-bounded algorithm for distributed optimization. In *Proc. of IJCAI*, 1452–1457.
- PETCU, A., FALTINGS, B. AND MAILLER, R. 2007. PC-DPOP: A new partial centralization algorithm for distributed optimization. In *Proc. of IJCAI*, 167–172.
- PETCU, A., FALTINGS, B. AND PARKES, D. 2008. M-DPOP: Faithful distributed implementation of efficient social choice problems. *Journal of Artificial Intelligence Research* 32, 705–755.
- PONTELLI, E., SON, T. C., BARAL, C. AND GELFOND, G. 2010. Logic programming for finding models in the logics of knowledge and its applications: A case study. *Theory and Practice of Logic Programming* 10, 4–6, 675–690.
- SADRI, F. AND TONI, F. 2003. Abductive logic programming for communication and negotiation amongst agents. ALP Newsletter.
- SAKAMA, C., SON, T. C. AND PONTELLI, E. 2011. A logical formulation for negotiation among dishonest agents. In *Proc. of IJCAI*, 1069–1074.
- SON, T. C., PONTELLI, E. AND SAKAMA, C. 2009. Logic programming for multiagent planning with negotiation. In *Proc. of ICLP*, 99–114.
- SULTANIK, E., LASS, R. AND REGLI, W. 2007. DCOPolis: A framework for simulating and deploying distributed constraint reasoning algorithms. In *Proc. of Distributed Constraint Reasoning Workshop*.
- UEDA, S., IWASAKI, A. AND YOKOO, M. 2010. Coalition structure generation based on distributed constraint optimization. In *Proc. of AAAI*, 197–203.
- VINYALS, M., RODRÍGUEZ-AGUILAR, J. AND CERQUIDES, J. 2011. Constructing a unifying theory of dynamic programming DCOP algorithms via the generalized distributive law. *Autonomous Agents and Multi-Agent Systems* 22, 3, 439–464.
- VLAHAVAS, I. 2002. MACLP: Multi agent constraint logic programming. *Information Sciences* 144, 1–4, 127–142.
- YEOH, W., FELNER, A. AND KOENIG, S. 2010. BnB-ADOPT: An asynchronous branch-and-bound DCOP algorithm. *Journal of Artificial Intelligence Research* 38, 85–133.
- YEOH, W., VARAKANTHAM, P. AND KOENIG, S. 2009. Caching schemes for DCOP search algorithms. In *Proc. of AAMAS*, 609–616.
- YEOH, W. AND YOKOO, M. 2012. Distributed problem solving. *AI Magazine* 33, 3, 53–65.
- ZIVAN, R., OKAMOTO, S. AND PELED, H. 2014. Explorative anytime local search for distributed constraint optimization. *Artificial Intelligence* 212, 1–26.