

PT-ISABB: A Hybrid Tree-based Complete Algorithm to Solve Asymmetric Distributed Constraint Optimization Problems

Yanchen Deng
College of Computer Science,
Chongqing University
Chongqing, China
dyc941126@126.com

Ziyu Chen*
College of Computer Science,
Chongqing University
Chongqing, China
chenziyu@cqu.edu.cn

Dingding Chen
College of Computer Science,
Chongqing University
Chongqing, China
dingding@cqu.edu.cn

Xingqiong Jiang
College of Computer Science,
Chongqing University
Chongqing, China
jxq@cqu.edu.cn

Qiang Li
College of Electrical Engineering,
Chongqing University
Chongqing, China
qiangli.ac@gmail.com

ABSTRACT

Asymmetric Distributed Constraint Optimization Problems (ADCOPs) have emerged as an important formalism in multi-agent community due to their ability to capture personal preferences. However, the existing search-based complete algorithms for ADCOPs can only use local knowledge to compute lower bounds, which leads to inefficient pruning and prohibits them from solving large scale problems. On the other hand, inference-based complete algorithms (e.g., DPOP) for Distributed Constraint Optimization Problems (DCOPs) require only a linear number of messages, but they cannot be directly applied into ADCOPs due to a privacy concern. Therefore, in the paper, we consider the possibility of combining inference and search to effectively solve ADCOPs at an acceptable loss of privacy. Specifically, we propose a hybrid complete algorithm called PT-ISABB which uses a tailored inference algorithm to provide tight lower bounds and a tree-based complete search algorithm to exhaust the search space. We prove the correctness of our algorithm and the experimental results demonstrate its superiority over other state-of-the-art complete algorithms.

KEYWORDS

ADCOP; Complete algorithms; Search; Inference

ACM Reference Format:

Yanchen Deng, Ziyu Chen, Dingding Chen, Xingqiong Jiang, and Qiang Li. 2019. PT-ISABB: A Hybrid Tree-based Complete Algorithm to Solve Asymmetric Distributed Constraint Optimization Problems. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*, IFAAMAS, 9 pages.

1 INTRODUCTION

Distributed Constraint Optimization Problems (DCOPs) [28] are a fundamental framework in multi-agent systems where agents cooperate with each other to optimize a global objective. DCOPs have

been successfully deployed in many real world applications including smart grids [6], radio frequency allocation [15], task scheduling [25], etc.

Algorithms for DCOPs can generally be classified into two categories, i.e., complete algorithms and incomplete algorithms. Search-based complete algorithms like SBB [10], AFB [7], ConFB [16], ADOPT [14] and its variants [9, 27] perform distributed searches to exhaust the search space, while inference-based complete algorithms including Action-GDL [26], DPOP [20] and its variants [21, 22] use dynamic programming to optimally solve problems. In contrast, incomplete algorithms including local search [12, 17, 29], GDL-based algorithms [3, 4, 24, 30] and sampling-based algorithms [5, 18] trade optimality for small computational efforts.

Asymmetric Distributed Constraint Optimization Problems (ADCOPs) [8] are a notable extension to DCOPs, which can capture ubiquitous asymmetric structures in real world scenarios [2, 13, 23]. That is, a constraint in an ADCOP explicitly defines the exact payoff for each participant instead of assuming equal payoffs for constrained agents. Solving ADCOPs is more challenging since algorithms must evaluate and aggregate the payoff for each participant of a constraint. ATWB and SABB [8] are asymmetric versions of AFB and SBB based on an one-phase strategy in which the algorithms systematically check each side of the constraints before reaching a full assignment. Besides, AsymPT-FB [11] is another search-based complete algorithm for ADCOPs, which implements a variation of forward bounding on a pseudo tree. However, to the best of our knowledge, there is no asymmetric adaptation of inference-based complete algorithms for DCOPs (e.g., DPOP). That is partially because these algorithms require the total knowledge of each constraint to perform variable elimination optimally. In other words, parent agents must surrender their private constraints to eliminate their children variables, which is unacceptable in a asymmetric scenario.

In this paper, we investigate the possibility of combining both inference and search to efficiently solve ADCOPs at an acceptable loss of privacy. Specifically, our main contributions are listed as follows.

*Corresponding author.

- We propose a hybrid tree-based complete algorithm for ADCOPs, called PT-ISABB.¹ The algorithm first uses a tailored version of ADPOP [19] to solve a subset of constraints, and the inference results stored in agents are served as look-up tables for tight lower bounds. Then, a variant of SABB is implemented on a pseudo tree to guarantee optimality.
- We theoretically show the completeness of our proposed algorithm. Moreover, we also prove that the lower bounds in PT-ISABB are at least as tight as the ones in AsymPT-FB when its maximal dimension limit $k = \infty$.
- We empirically evaluate our algorithm on various benchmarks. Our study shows that PT-ISABB requires significantly fewer messages and lower NCLOs than state-of-the-art search-based complete algorithms including AsymPT-FB. The experimental results also indicate that our proposed algorithm leaks less privacy than AsymPT-FB when solving complex problems.

2 BACKGROUND

In this section, we review the preliminaries including ADCOPs, pseudo tree, DPOP and ADPOP.

2.1 Asymmetric Distributed Constraint Optimization Problems

An asymmetric distributed constraint optimization problem can be defined by a tuple $\langle A, X, D, F \rangle$ in which

- $A = \{a_1, \dots, a_q\}$ is a set of agents
- $X = \{x_1, \dots, x_n\}$ is a set of variables
- $D = \{D_1, \dots, D_n\}$ is a set of finite and discrete domains. Each variable x_i takes a value from D_i
- $F = \{f_1, \dots, f_m\}$ is a set of constraints. Each constraint $f_i : D_{i_1} \times \dots \times D_{i_k} \rightarrow \mathbb{R}_+^k$ defines a set of non-negative costs for every possible value combination of the set of variables it is involved in

For the sake of simplicity, we assume that an agent only controls a variable and all constraints are binary. Therefore, the term *agent* and *variable* can be used interchangeably. Besides, for the constraint between x_i and x_j , we denote the private cost functions for x_i and x_j as f_{ij} and f_{ji} , respectively. Note that in the asymmetric setting, f_{ij} does not necessarily equal to f_{ji} . A solution to an ADCOP is the assignments to all the variables with the minimal aggregated cost. An ADCOP can be visualized by a *constraint graph* in which the vertexes denote the variables and the edges denote the constraints between agents. Fig. 1 (a) visualizes an ADCOP with four agents and four constraints.

2.2 Pseudo Tree

A pseudo tree is an ordered arrangement to a constraint graph in which agents in different branches are independent and thus search can be performed in parallel on these independent branches. A pseudo tree can be generated by a depth-first traverse of the constraint graph, which categorizes the constraints into tree edges and pseudo edges. For an agent a_i , we denote its parent as $P(a_i)$ which is the ancestor connecting to a_i through a tree edge, its




Figure 1: An example of constraint graph and pseudo tree

pseudo parents as $PP(a_i)$ which is a set of ancestors connecting to a_i through pseudo edges, its children and pseudo children as $C(a_i)$ and $PC(a_i)$ which are the sets of descendants connecting to a_i via tree edges and pseudo edges, respectively. For the sake of clarity, we denote all the parents of a_i as $AP(a_i) = PP(a_i) \cup \{P(a_i)\}$. We also denote its separator, i.e., the set of ancestors that are constrained with a_i or its descendants, as $Sep(a_i)$ [21]. Fig. 1 (b) gives a possible pseudo tree deriving from Fig. 1 (a).

2.3 DPOP and ADPOP

DPOP is an important inference-based algorithm that performs dynamic programming on a pseudo tree, starting with a phase of utility propagation. In the phase, each agent joins the received utilities from its children with its local utility, eliminates its dimension by calculating the optimal utility for each assignment combination of its separator, and propagates the reduced utility to its parent. After that, a value propagation phase starts from the root agent. In the phase, each agent chooses the optimal assignment according to the utilities calculated in the previous phase and the assignments from its parent, and broadcasts the extended assignments to its children. The algorithm terminates when all agents have chosen their optimal assignments.

Although DPOP only requires a linear number of messages to solve a DCOP, its memory consumption is still exponential in induced width, which prohibits it from solving more complex problems. Thus, Petcu et. al, proposed ADPOP which is an approximate version of DPOP and allows the desired trade-off between solution quality and computational complexity. Specifically, ADPOP imposes a limit $maxDim$ on the maximum number of dimensions in each message. When the number of dimensions in an outgoing message exceeds the limit, the algorithm drops a set of dimensions to stay below the limit. That is, the algorithm computes an upper bound and a lower bound by applying a maximal/minimal projection on these dimensions. During the value propagation phase, agents can make decisions according to the highest utilities in either upper bounds or lower bounds.

3 PROPOSED METHOD

In this section, we present our proposed PT-ISABB, a two-phase hybrid complete algorithm for ADCOPs. We begin with a motivation, and then present the details of the inference phase and the search phase, respectively.

¹The source code is available at https://github.com/czy920/DCOPsolvlerAlgorithm_PTISABB.

Algorithm 1: Inference phase for a_i

```

When Initialization:
1   $join_i \leftarrow copy(local\_util_i)$ 
2  if  $a_i$  is a leaf then
3  | SendUtil()
When received UTIL ( $child\_util$ ) from  $a_c \in C(a_i)$ :
4   $Child\_util_i^c \leftarrow \min_{x_c} child\_util \otimes f_{ic}$ 
5   $join_i \leftarrow join_i \otimes Child\_util_i^c$ 
6  if  $a_i$  has received all UTIL from  $C(a_i)$  then
7  | if  $a_i$  is the root then
8  | | start Search phase
9  | else
10 | | SendUtil()
Function SendUtil():
11 if  $|dim(join_i)| > k$  then
12 | select  $S \subset dim(join_i) \setminus x_i$ , s.t.  $|S| = |dim(join_i)| - k$ 
13 |  $join_i \leftarrow \min_S join_i$ 
14 | send UTIL( $join_i$ ) to  $P(a_i)$ 

```

Figure 2: Pseudo code of inference phase

3.1 Motivation

The existing search-based complete algorithms for ADCOPs can only use *local* knowledge to compute lower bounds, which leads to inefficient pruning. More specifically, unassigned agents report the best *local* costs under the given partial assignments to compute lower bounds. Taking Fig. 1 as an example, in AsymPT-FB agent a_1 can receive **LB_reports** from a_2 and a_3 . As a consequence, a_1 can only be aware of the lower bounds of f_{21} and f_{31} and does not have any knowledge about the remaining constraints (i.e., the constraints between a_2 and a_3 , a_2 and a_4). On the other hand, inference algorithms like DPOP are able to aggregate and propagate the *global* utility, but they are not applicable to ADCOPs due to a privacy concern. For example, a_3 needs to know both f_{13} and f_{23} to optimally eliminate x_3 , which violates the privacies of a_1 and a_2 . Thus, to overcome the pathologies, we propose a novel hybrid scheme to solve ADCOPs, which combines both inference and search. Specifically, the scheme consists of the following phases.

- **Inference phase:** performing a bottom-up utility propagation with respect to a subset of constraints to build look-up tables for lower bounds
- **Search phase:** using a tree-based complete search algorithm for ADCOPs to exhaust the search space and guarantee optimality

In this paper, we propose a tailored version of ADPOP for the inference phase to avoid the severe privacy loss and exponential memory consumption in DPOP. Furthermore, we implement SABB on a pseudo tree for the search phase and propose an algorithm called PT-ISABB. Although they both operate on pseudo trees, our algorithm excels AsymPT-FB twofold. When its maximal dimension limit $k = \infty$, the lower bounds in our algorithm are at least as tight as the ones in AsymPT-FB (see Property 4.1 for detail). Moreover, PT-ISABB avoids to perform forward bounding which is expensive during the search phase.

3.2 Inference Phase

Fig. 2 gives the sketch of the inference phase for PT-ISABB. The phase begins with leaf agents who send their local utilities to their parents via UTIL messages (line 1 - 3). Particularly, if the number of dimensions in the utility exceeds the limit k (line 11), we drop the dimensions of the highest ancestors to stay below the limit by

a minimal projection (line 12 - 13). Here, $local_util_i$ denotes the combination of the constraints between agent a_i and its parent and pseudo parents enforced in a_i side, i.e.,

$$local_util_i = \bigotimes_{a_j \in AP(a_i)} f_{ij}$$

Note that in our algorithm we do not require the parent agents to disclose their private functions to perform inference exactly. In this way, a local utility table only involves the functions of that agent and the privacies of its parent and pseudo parents are therefore guaranteed. On the other hand, however, ignoring the private functions of parents and pseudo parents leads to severe inconsistencies when performing variable elimination. In other words, we actually trade lower bound tightness for privacy. We try to alleviate the problem by performing *non-local* elimination which is elaborated as follows.

When a_i receives a UTIL message from its child a_c , it joins the utility from a_c with its corresponding private constraint function and then eliminates the dimension x_c for a more complete utility $Child_util_i^c$ (line 4). Compared to DPOP and ADPOP, the elimination of each variable is postponed to its parent in the pseudo-tree. Taking Fig. 1 (b) for example, the UTIL message from a_3 to a_2 is given by

$$f_{32} + f_{31}$$

if the maximal dimension limit $k \geq 3$, and the elimination of x_3 is actually performed by a_2 . That is,

$$Child_util_2^3 = \min_{x_3} (f_{23} + f_{32} + f_{31})$$

Then, a_i initiates the search phase after receiving all UTIL messages from its children if it is the root agent (line 6 - 8). Otherwise, it propagates the joint utility to its parent (line 9 - 10).

3.3 Search Phase

The phase performs a branch-and-bound search on a pseudo tree to exhaust the search space. Specifically, each branching agent decomposes the problem into several subproblems, and each of its children solves a subproblem in parallel. To detect and discard the suboptimal solution, each agent maintains an upper bound for its subproblem and a lower bound for each value in its domain. Therefore, each agent a_i needs to maintain the following data structures.

- $Srch_val_i^c$ records the assignment currently being explored in the subtree rooted at $a_c \in C(a_i)$. The data structure is necessary because asynchronous search is carried out in parallel in sub-trees based on different possible values of x_i .
- $high_cost_i(d_i)$ is the cost for $d_i \in D_i$ between a_i and its parent and pseudo parents under the current partial assignment (Cpa), which is initially set to the cost enforced in a_i side. That is,

$$high_cost_i(d_i) = \sum_{a_j \in AP(a_i)} f_{ij}(d_i, Cpa(x_j)) \quad (1)$$

- $lb_i^c(d_i)$ is the lower bound of child $a_c \in C(a_i)$ for $d_i \in D_i$, which is initially set to the utility under Cpa and d_i . That is,

$$lb_i^c(d_i) = Child_util_i^c(Cpa_{[Sep(a_c)]}, x_i = d_i) \quad (2)$$

where $Cpa_{[Sep(a_c)]}$ is a slice to Cpa under $Sep(a_c)$, i.e.,

$$Cpa_{[Sep(a_c)]} = \{(x_j, d_j) | (x_j, d_j) \in Cpa \wedge x_j \in Sep(a_c)\}$$

Algorithm 2: Search phase for a_i (message passing)

```

When Initialization:
1  if  $a_i$  is the root then
2    InitializeVariables()
3     $d_i \leftarrow$  the first element in  $D_i$ ,  $Cpa \leftarrow (x_i, d_i)$ 
4    send CPA( $Cpa, \infty$ ) to  $a_c$ ,  $\forall a_c \in C(a_i)$ 
When received CPA( $Cpa, ub_i$ ) from  $P(a_i)$ :
5    store  $\{Cpa, ub_i\}$ 
6    InitializeVariables()
7     $d_i \leftarrow$  NextFeasibleAssignment( $null$ )
8    if  $d_i \neq null$  then
9      send COST_REQ( $Cpa(x_j), d_i$ ) to  $a_j$ ,  $\forall a_j \in AP(a_i)$ 
10      $Srch\_val_i^c \leftarrow d_i$ ,  $\forall a_c \in C(a_i)$ 
11    else
12     SendBacktrack()
When received COST_REQ( $d_j, d_c$ ) from  $a_c \in C(a_i) \cup PC(a_i)$ :
13    send COST( $f_{ic}(d_i, d_c), d_c$ ) to  $a_c$ 
When received COST( $cost, d_j$ ) from  $a_p \in AP(a_i)$ :
14     $high\_cost_i(d_i) \leftarrow high\_cost_i(d_i) + cost$ 
15    if  $a_i$  has not received all COST from  $AP(a_i)$  for  $d_i$  then
16     return
17    if  $a_i$  is a leaf then
18      $ub_i \leftarrow \min(ub_i, lb_i(d_i))$ 
19      $d'_i \leftarrow$  NextFeasibleAssignment( $d_i$ )
20     if  $d'_i \neq null$  then
21       send COST_REQ( $Cpa(x_j), d'_i$ ) to  $a_j$ ,  $\forall a_j \in AP(a_i)$ 
22     else
23      SendBacktrack()
24    else
25     if  $lb_i(d_i) < ub_i$  then
26       foreach  $a_c \in C(a_i) \wedge Srch\_val_i^c = d_i$  do
27         SendCpa( $d_i, a_c$ )
28     else
29        $d'_i \leftarrow$  NextFeasibleAssignment( $d_i$ )
30        $Srch\_val_i^c \leftarrow d'_i$ ,  $\forall a_c \in C(a_i) \wedge Srch\_val_i^c = d_i$ 
31       if  $d'_i \neq null$  then
32         send COST_REQ( $Cpa(x_j), d'_i$ ) to  $a_j$ ,  $\forall a_j \in AP(a_i)$ 
33       else if  $\forall a_c \in C(a_i)$ ,  $Srch\_val_i^c = null$  then
34         SendBacktrack()
When received BACKTRACK( $lb_i^*, Spa_i$ ) from  $a_c \in C(a_i)$ :
35     $d_i \leftarrow Srch\_val_i^c$ ,  $lb_i^*(d_i) \leftarrow lb_i^*$ ,  $Spa_i(d_i) \leftarrow Spa_i(d_i) \cup Spa_c$ 
36    if  $a_i$  has received all BACKTRACK from  $C(a_i)$  for  $d_i$  then
37      $Cmplt\_vals_i \leftarrow Cmplt\_vals_i \cup \{d_i\}$  //marks  $d_i$  as complete
38      $ub_i \leftarrow \min(ub_i, lb_i(d_i))$ 
39      $d'_i \leftarrow$  NextFeasibleAssignment( $d_i$ ),  $Srch\_val_i^c \leftarrow d'_i$ 
40     if  $d'_i \neq null$  then
41       if  $a_i$  has received all COST from  $AP(a_i)$  for  $d'_i$  then
42         SendCpa( $d'_i, a_c$ )
43       else if  $a_i$  has not requested costs for  $d'_i$  then
44         send COST_REQ( $Cpa(x_j), d'_i$ ) to  $a_j$ ,  $\forall a_j \in AP(a_i)$ 
45     else if  $\forall a_c \in C(a_i)$ ,  $Srch\_val_i^c = null$  then
46       if  $a_i$  is the root then
47         send TERMINATE to  $a_c$ ,  $\forall a_c \in C(a_i)$ 
48         terminate
49       else
50        SendBacktrack()
When received TERMINATE from  $P(a_i)$ :
51    send TERMINATE to  $a_c$ ,  $\forall a_c \in C(a_i)$ 
52    terminate

```

Figure 3: Pseudo code of search phase (message passing)

When a_i receives a BACKTRACK message from a_c , it replaces the initial lower bound with the actual cost reported by a_c (or ∞ if d_i is infeasible for a_c given Cpa).

- $lb_i(d_i)$ is the lower bound for $d_i \in D_i$, i.e.,

$$lb_i(d_i) = high_cost_i(d_i) + \sum_{a_c \in C(a_i)} lb_i^c(d_i) \quad (3)$$

- $Cmplt_vals_i$ is the set of assignments for which a_i has received all BACKTRACK messages from its children, and is initially set to \emptyset . Particularly, $Cmplt_vals_i = D_i$ if a_i is a leaf agent.
- lb_i^* is the best cost explored under Cpa , which is given by

$$lb_i^* = \min_{d_i \in Cmplt_vals_i} lb_i(d_i) \quad (4)$$

Particularly, if $Cmplt_vals_i = \emptyset$, $lb_i^* = \infty$.

Algorithm 2: Search phase for a_i (auxiliary functions)

```

Function InitializeVariables():
53  initialize  $Cmplt\_vals_i$ 
54  initialize  $high\_cost_i(d_i)$ ,  $Spa_i(d_i)$ ,  $\forall d_i \in D_i$ 
55  initialize  $Srch\_val_i^c$ ,  $\forall a_c \in C(a_i)$ 
56  initialize  $lb_i^c(d_i)$ ,  $\forall a_c \in C(a_i)$ ,  $d_i \in D_i$ 
Function NextFeasibleAssignment( $d_i$ ):
57   $d'_i \leftarrow$  the element next to  $d_i$  in  $D_i$ 
58  while  $d'_i \neq null \wedge lb_i(d'_i) \geq ub_i$  do
59     $d'_i \leftarrow$  the next element in  $D_i$ 
60  return  $d'_i$ 
Function SendCpa( $d_i, a_c$ ):
61   $TmpCpa \leftarrow Cpa \cup (x_i, d_i)$ 
62   $ub_c \leftarrow ub_i - high\_cost_i(d_i) - \sum_{a_j \in C(a_i) \wedge j \neq c} lb_i^j(d_i)$ 
63  send CPA( $TmpCpa, ub_c$ ) to  $a_c$ 
Function SendBacktrack():
64  if  $Cmplt\_vals_i = \emptyset$  then
65    send BACKTRACK( $\infty, \emptyset$ ) to  $P(a_i)$ 
66  else
67     $d_i^* \leftarrow \arg \min_{d_i \in Cmplt\_vals_i} lb_i(d_i)$ 
68    send BACKTRACK( $lb_i^*, Spa_i(d_i^*)$ ) to  $P(a_i)$ 

```

Figure 4: Pseudo code of search phase (auxiliary functions)

- $Spa_i(d_i)$ is the optimal assignment to its subproblem under Cpa when $x_i = d_i$ and is initially set to $\{x_i = d_i\}$. Particularly, $Spa_i(d_i^*)$ is the optimal solution if a_i is the root agent, where $d_i^* = \arg \min_{d_i \in Cmplt_vals_i} lb_i(d_i)$.

Fig. 3 and Fig. 4 give the pseudo codes of the search phase for PT-ISABB. The phase begins with the root agent sending the first element in its domain to its children (line 1 - 4). When an agent a_i receives a CPA message from its parent, it first stores the partial assignment Cpa and upper bound ub_i and then finds the first feasible assignment d_i such that $lb_i(d_i) < ub_i$ (line 57 - 60). If such an assignment exists, a_i sends COST_REQ messages to its parent and pseudo parents to request the private costs of other side for d_i (line 8 - 10, line 13). Otherwise, it sends a BACKTRACK message with an infinity cost and an empty subproblem assignment (line 64 - 65) to its parent to announce that the given Cpa is infeasible (line 11 - 12).

When a_i receives a COST message for d_i , it adds the other side cost to $high_cost_i(d_i)$ (line 14). After receiving all the COST messages for d_i from its parent and pseudo parents, a_i is able to determine whether it should continue to explore d_i . If a_i is a leaf agent, it just updates the current upper bound and switches to the next feasible assignment d'_i (line 17 - 19) since the search space no longer needs to be expanded. If such d'_i exists, a_i requests costs for the assignment (line 20 - 21). Otherwise, it backtracks to its parent by reporting the best cost and the best subproblem assignment explored under Cpa (line 22 - 23, line 66 - 68). If a_i is not a leaf agent and the current lower bound for d_i is still less than its upper bound, it expands the search space by sending CPA messages to its children who are going to explore d_i (line 25 - 27). Each message contains an extended partial assignment (line 61) and an upper bound which is the remainder after deducting the cost incurred by d_i and the lower bounds of the other children from a_i 's upper bound (line 62). Otherwise, d_i is proven to be suboptimal and the agent switches to the next feasible assignment (line 28 - 30). If such an assignment exists, a_i requests costs for it (line 31 - 32). A backtrack takes place if all children exhaust a_i 's domain (line 33 - 34).

When a_i receives a BACKTRACK message for d_i from a child a_c , it updates the corresponding lower bound $lb_i^c(d_i)$ with the actual cost lb_c^* reported by a_c if Cpa and the assignment d_i is feasible (otherwise $lb_c^* = \infty$), and merges the best assignments from a_c (line 35). If a_i has received all the BACKTRACK messages for d_i from its children, it marks d_i as complete and updates the current upper bound for its subproblem (line 36 - 38). a_i also needs to determine the next assignment d'_i for a_c to explore (line 39). If d'_i exists and a_i has received all the COST messages from its parent and pseudo parents, it informs a_c to explore d'_i by sending a CPA message (line 40 - 42). Otherwise, a_i requests costs for d'_i if it has not been done (line 43 - 44). If d'_i does not exist and all children have exhausted a_i 's domain, a_i informs its children to terminate and terminates itself if it is the root agent (line 45 - 48). Otherwise, it backtracks to its parent (line 49 - 50).

4 THEORETICAL RESULTS

4.1 Correctness

In this section, we first prove the termination and optimality, and further establish the completeness of PT-ISABB.

LEMMA 4.1. *PT-ISABB will terminate after a finite number of iterations.*

PROOF. Directly from the pseudo codes, the inference phase will terminate since it only needs a linear number of messages. Thus, to prove the termination, it is enough to show that the same partial assignment cannot be explored twice in the search phase, i.e., an agent will not receive two identical $Cpas$. Obviously, the claim holds for the root agent since it does not receive any CPA message. For an agent and a given Cpa from its parent, it will send several CPA messages to each child. Since each of them contains the different assignments of the agent (line 29, line 39, line 57 - 60), the $Cpas$ sent to the child are all different. Therefore, the termination is hereby guaranteed. \square

LEMMA 4.2. *For an agent a_i and a given Cpa , the cost incurred by any assignment Spa_i to the subtree rooted at a_i with the assignment ($x_i = d_i$) is no less than the corresponding lower bound $lb_i^c(d_i)$.*

PROOF. The lemma for a leaf is trivial since $lb_i(d_i)$ is set to the cost of d_i which is obviously no greater than the cost of the feasible assignment. We now focus on no-leaf agents. Recall that a_i will replace the original lower bound with the actual cost reported by a_c after receiving a BACKTRACK message for d_i from $a_c \in C(a_i)$ (line 35). Thus, to prove the lemma, it is sufficient to show that the initial lower bound $lb_i^c(d_i)$ is no greater than the actual cost of $Spa_c, \forall a_c \in C(a_i)$, where $Spa_c \subset Spa_i$ is the assignment to the subtree rooted at a_c .

Consider the induction basis, i.e., a_i 's children are leaves. For each child $a_c \in C(a_i)$, we have

$$\begin{aligned} cost(Spa_c) &= \sum_{a_j \in AP(a_c)} f_{jc}(d_j, d_c) + f_{cj}(d_c, d_j) \\ &\geq \min_{x_c} f_{ic}(d_i, x_c) + \sum_{a_j \in AP(a_c)} f_{cj}(x_c, d_j) \\ &\geq Child_util_i^c(Cpa_{[Sep(a_c)]}, x_i = d_i) = lb_i^c(d_i) \end{aligned}$$

where d_i is the assignment to x_i in Cpa or Spa_c . The equation in the second to the last step holds when the maximal dimension limit $k = \infty$. Thus, the lemma holds for the basis.

Assume that the lemma holds for all $a_c \in C(a_i)$. Next, we are going to show the lemma holds for a_i as well. For each child $a_c \in C(a_i)$, we have

$$\begin{aligned} cost(Spa_c) &= \sum_{a_j \in AP(a_c)} f_{cj}(d_c, d_j) + f_{jc}(d_j, d_c) + \sum_{a_{c'} \in C(a_c)} cost(Spa_{c'}) \\ &\geq \sum_{a_j \in AP(a_c)} f_{cj}(d_c, d_j) + f_{jc}(d_j, d_c) + \sum_{a_{c'} \in C(a_c)} lb_{c'}^c(d_c) \\ &\geq \min_{x_c} f_{ic}(d_i, x_c) + \sum_{a_j \in AP(a_c)} f_{cj}(x_c, d_j) + \sum_{a_{c'} \in C(a_c)} lb_{c'}^c(x_c) \\ &\geq Child_util_i^c(Cpa_{[Sep(a_c)]}, x_i = d_i) = lb_i^c(d_i) \end{aligned}$$

which establishes the lemma. \square

LEMMA 4.3. *For an agent a_i and a given Cpa , any assignment to the subtree rooted at a_i with cost greater than ub_i cannot be a part of a solution with cost less than the global upper bound.*

PROOF. We will prove recursively, by showing that for a partial assignment Spa_i to the subtree rooted at a_i with $cost(Spa_i) > ub_i$, any partial assignment $Spa_j \supset Spa_i$ to the subtree rooted at a_j will have $cost(Spa_j) > ub_j$ where $a_j = P(a_i)$. Note that ub_j could be either an upper bound from a_j via a CPA message (line 5) or a result of updating the upper bound locally (line 18, line 38). a_i cannot backtrack by reporting Spa_i in the latter case since there must exist a better partial assignment whose cost is ub_i . If ub_i is received from a_j , according to line 62, we have

$$ub_j = ub_i + high_cost_j(d_j) + \sum_{a_c \in C(a_j) \wedge c \neq i} lb_j^c(d_j)$$

Thus, $cost(Spa_i) > ub_i$ necessarily means that any partial assignment $Spa_j \supset Spa_i$ will have $cost(Spa_j) > ub_j$. \square

THEOREM 4.4. *PT-ISABB is complete.*

PROOF. Immediately from Lemma 4.1, Lemma 4.2 and Lemma 4.3, the algorithm will terminate and all pruned assignments are suboptimal. Thus, PT-ISABB is complete. \square

4.2 Lower bound tightness

PROPERTY 4.1. *For an agent a_i and a given Cpa , the initial lower bound $lb_i^c(d_i)$ of $a_c \in C(a_i)$ for d_i is at least as tight as the one in AsymPT-FB when the maximal dimension limit $k = \infty$.*

PROOF. In AsymPT-FB, the lower bound for a_c after receiving all the **LB_Reports** from the subtree rooted at a_c is given by the sum of the best single side local costs of a_c 's descendants under Cpa . That is,

$$\begin{aligned} SubtreeLB_i^c(d_i) &= \sum_{a_j \in Desc(a_c)} \min_{x_j} \sum_{a_l \in Sep(a_c) \cap PPP(a_j)} f_{jl}(x_j, d_l) \\ &\quad + \min_{x_c} \sum_{a_l \in AP(a_c)} f_{cl}(x_c, d_l) \end{aligned}$$

where $Desc(a_c)$ is the set of the descendants of a_c . For the sake of clarity, we denote the vector of x_c and its descendant variables as

\mathbf{x}_c . Next, we will show $lb_i^c(d_i) \geq \text{SubtreeLB}_i^c(d_i)$. Since $k = \infty$, the inference phase does not drop any dimension. Thus, we have

$$\begin{aligned} lb_i^c(d_i) &= \min_{\mathbf{x}_c} \sum_{a_j \in Desc(a_c)} \left(\sum_{a_l \in AP(a_j) \cap Sep(a_c)} f_{jl}(x_j, d_l) \right. \\ &\quad \left. + \sum_{a_l \in AP(a_j) \cap Desc(a_i)} f_{jl}(x_j, x_l) + \sum_{a_l \in C(a_j)} f_{jl}(x_j, x_l) \right) \\ &\quad + \sum_{a_l \in AP(a_c)} f_{cl}(x_c, d_l) + \sum_{a_l \in C(a_c)} f_{cl}(x_c, x_l) + f_{ic}(d_i, x_c) \\ &\geq \min_{\mathbf{x}_c} \sum_{a_j \in Desc(a_c)} \sum_{a_l \in AP(a_j) \cap Sep(a_c)} f_{jl}(x_j, d_l) \\ &\quad + \sum_{a_l \in AP(a_c)} f_{cl}(x_c, d_l) \end{aligned}$$

Since $x_c \in \mathbf{x}_c$ and $AP(a_j) \supset PP(a_j)$, the right-hand side of the inequality in the last step can be further reduced. That is,

$$\begin{aligned} lb_i^c(d_i) &\geq \min_{\mathbf{x}_c \setminus x_c} \sum_{a_j \in Desc(a_c)} \sum_{a_l \in AP(a_j) \cap Sep(a_c)} f_{jl}(x_j, d_l) \\ &\quad + \min_{\mathbf{x}_c} \sum_{a_l \in AP(a_c)} f_{cl}(x_c, d_l) \\ &\geq \min_{\mathbf{x}_c \setminus x_c} \sum_{a_j \in Desc(a_c)} \sum_{a_l \in PP(a_j) \cap Sep(a_c)} f_{jl}(x_j, d_l) \\ &\quad + \min_{\mathbf{x}_c} \sum_{a_l \in AP(a_c)} f_{cl}(x_c, d_l) \\ &\geq \sum_{a_j \in Desc(a_c)} \min_{x_j} \sum_{a_l \in Sep(a_c) \cap PP(a_j)} f_{jl}(x_j, d_l) \\ &\quad + \min_{\mathbf{x}_c} \sum_{a_l \in AP(a_c)} f_{cl}(x_c, d_l) \\ &= \text{SubtreeLB}_i^c(d_i) \end{aligned}$$

which concludes the property. \square

4.3 Complexity

Since an agent a_i stores $Child_util_i^c$ and $lb_i^c(d_i)$ for each child, the overall space complexity in the worst case (i.e., $k = \infty$) is $O(|C(a_i)|d_{max}^{|Sep(a_i)|+1} + |C(a_i)||D_i|)$ where $d_{max} = \max_{a_j \in Sep(a_i)} |D_j|$. Since it contains all the dimensions of $Sep(a_i)$ and itself, the size of an UTIL message from a_i is $O(d_{max}^{|Sep(a_i)|+1})$ when $k = \infty$. For a CPA message, it consists of the assignment of each agent and an upper bound. Thus, the size of a CPA message is $O(|A|)$. Other messages including COST_REQ, COST, BACKTRACK and TERMINATE carry several scalars and thus they only require $O(1)$ space.

Different than standard DPOP/ADPOP, PT-ISABB only requires $|A| - 1$ messages in the inference phase since it does not have the value propagation phase. Like any other search based complete algorithm, the message number of the search phase grows exponentially with respect to the agent number.

5 EXPERIMENTAL RESULTS

We empirically evaluate PT-ISABB with state-of-the-art search-based complete algorithms for ADCOPs including SABB, ATWB and AysmPT-FB on three configurations. To demonstrate the real power of non-local elimination, we also consider SABB on a pseudo-tree (PT-SABB) and the local elimination version of PT-ISABB (PT-ISABB, local) with $k = \infty$. In the first ADCOP configuration, we set the graph density to 0.25, the domain size to 3 and vary the agent number from 8 to 18. The second configuration is ADCOPs with 8

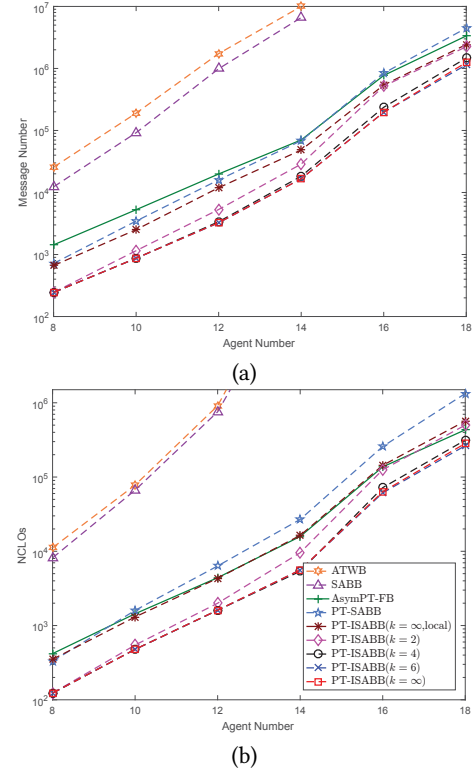


Figure 5: Performance comparison under different agent numbers

agents and the domain size of 8. The graph density varies from 0.25 to 1. In the last configuration, we consider asymmetric MaxDCSPs with 10 agent, the domain size of 10 and the graph density of 0.4, and the tightness varies from 0.1 to 0.8. For each of the settings, we generate 50 random instances and the results are averaged over all instances. In our experiments, we use the number of non-concurrent logical operations (NCLo) [16] to evaluate hardware-independent runtime, in which the logical operations in the inference phase are accesses to utility tables, and for the search phase and other competitors they are constraint checks. Also, we use the message number and the size of total information exchanged to measure the network load. For asymmetric MaxDCSPs, we use entropy [1] to quantify the privacy loss [8, 11]. The experiments are conducted on an i7-7820x workstation with 32GB of memory and for each algorithm we set the timeout to 2 minutes.

Fig. 5 gives the performance comparison on different agent numbers, and the average induced widths in the experiments are 1 ~ 6.84. It can be seen from the figure that all the algorithms suffer from exponential overheads as the agent number grows. Among them, our proposed PT-ISABB requires significant fewer messages and lower NCLOs than the other competitors, which demonstrates the superiority of the hybrid execution of inference and search. On the other hand, although PT-ISABB ($k = \infty$, local) employs an complete one-side inference to construct the initial lower bounds, it is still inferior to PT-ISABB with $k > 2$, which demonstrate the necessity of non-local elimination. Besides, it is worth noting that

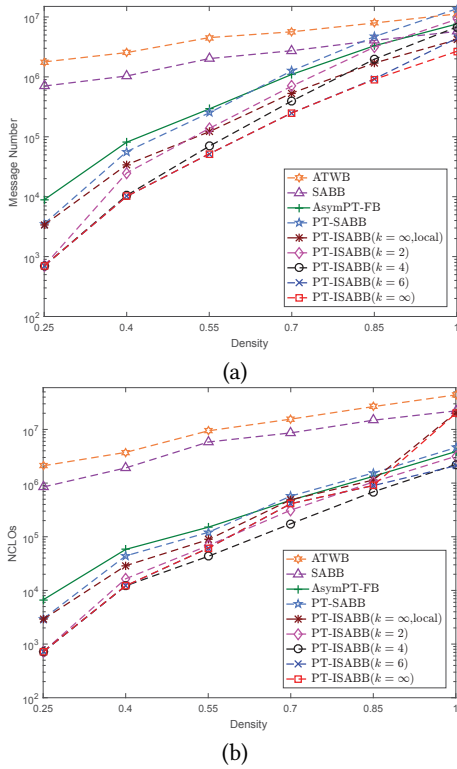


Figure 6: Performance comparison under different graph densities

PT-ISABB requires much fewer messages than AsymPT-FB even when the maximal dimension limit k is small (e.g., $k = 2$). That is because PT-ISABB does not rely on forward bounding which is expensive in message-passing to compute lower bounds. Moreover, the phenomenon also indicates that our algorithm can produce tighter lower bounds even if the memory budget is relatively low.

Fig. 6 gives the results under different graph densities. The average induced widths here are $1 \sim 6$. Note that in this configuration, the size of the search space does not change and the complexity is reflected in the topologies. It can be concluded from the figure that all the tree-based algorithms exhibit great superiorities when the graph density is low, and the advantages vanish as the density grows. That is because those algorithms can effectively parallel the search processes on sparse problems. Dense problems, on the other hand, usually result in pseudo trees with low branching factors, making the tree-based algorithms require more messages than SABB. Even so, our proposed PT-ISABB with large k still outperforms SABB when the problems are fully connected, which demonstrates the necessity of tighter lower bounds. Additionally, the figure also indicates that PT-ISABB with different k performs similarly on sparse problems, but the performances vary a lot on dense problems. That is due to the fact that the induced width of a pseudo tree is relatively small when solving a sparse problem and thus only a small set of dimensions is dropped during the inference phase. Besides, it can be seen from the figure that although both PT-ISABB ($k = \infty, \text{local}$) and PT-ISABB ($k = \infty$) perform complete

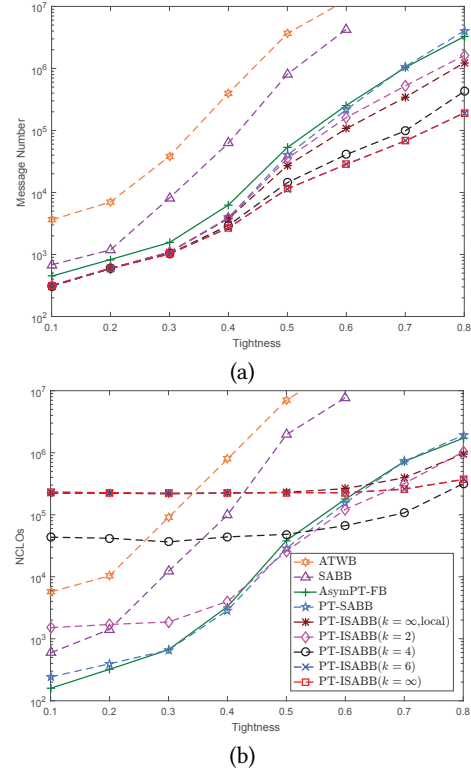


Figure 7: Performance comparison under different tightness

Algorithm	Density	0.25	0.4	0.55	0.7	0.85	1
ATWB		91,273.29	97,715.94	175,594.67	190,195.50	266,654.30	373,255.37
SABB		64,264.44	68,342.30	142,308.71	164,670.95	233,956.11	329,344.13
AsymPTFB		168.07	1,174.82	3,161.29	9,483.76	23,958.52	57,261.91
PT-SABB		85.37	990.77	3,260.30	12,920.10	35,221.34	103,409.30
PT-ISABB($k = \infty, \text{local}$)		67.41	581.41	2,137.57	8,653.52	25,444.45	73,442.96
PT-ISABB($k = 2$)		8.59	343.04	1,661.95	7,426.99	23,572.68	69,207.66
PT-ISABB($k = 4$)		8.58	143.45	811.75	3,812.25	14,755.51	50,844.86
PT-ISABB($k = 6$)		8.57	147.73	897.09	3,199.69	9,734.17	35,617.43
PT-ISABB($k = \infty$)		8.59	147.55	897.39	3,198.51	1,0435.57	94,536.77

Table 1: The size of total information exchanged of each algorithm under different graph densities (in KB)

inferences, the non-local elimination version requires lower NCLOs and fewer message numbers in most of the settings. That is because the non-local elimination version can provide tighter lower bounds which result in efficient pruning, and thus the algorithm incurs fewer constraint checks and messages in the search phase.

Table 1 presents the size of total information exchanged of each algorithm under different densities. It can be seen from the table that all the non-local elimination versions of PT-ISABB exhibit great advantages over the other search-based competitors, except that PT-ISABB ($k = 2$) is slightly inferior to AsymPT-FB when solving the fully-connected problems. The phenomenon indicates that although a message in the inference phase is generally larger than the one in the search phase, the algorithms can still gain great reductions on network traffics since the effective pruning in the search phase greatly reduces the message number. Besides, it is interesting to find that a large dimension limit (e.g., $k = \infty$) does

not necessarily result in the smallest traffic. That is because the size of a message in the inference phase is exponential to the minimum of the induced width and the dimension limit. Besides, it should be noted that although PT-SABB requires more messages than ATWB and SABB when solving fully-connected problems according to Fig. 6, it still incurs much smaller traffic due to the fact that the last agent in ATWB and SABB needs to broadcast the reached complete solution to all other agents once a new solution is constructed. In contrast, agents in PT-SABB only back up the assignments to their descendants, which are subsets of the complete solution, to their parents via BACKTRACK messages.

Fig. 7 presents the results on asymmetric MaxDCSPs with different tightness, and the average induced width is 3.92. This configuration neither increases the search space nor affects the topologies, but instead increases the difficulty of pruning. All the algorithms except ATWB produces few messages when solving problems with low tightness. That is because on these problems the algorithms can find low upper bounds very quickly to prune most of the search space. As the tightness grows, the number of prohibited combinations increases and the algorithms can no longer find low upper bounds promptly. As a result, the algorithms require much more search efforts to exhaust the search space. Since they cannot exploit topologies to accelerate the search process, SABB and ATWB perform poorly and can only solve the problems with tightness up to 0.6. On the other hand, the tree-based algorithms divide a problem to several smaller subproblems at each branching agent and search the subproblems in parallel. Thus, those algorithms exhibit better performances and solve all the problems. Among them, our proposed PT-ISABB with $k \geq 4$ incurs much smaller overheads, which demonstrates the effectiveness of the inference phase in computing tighter lower bounds. In other words, although PT-ISABB only guarantees to produce lower bounds as tight as the ones of AsymPT-FB when $k = \infty$ according to Property 4.1, it requires less memory consumption to compute such lower bounds in practice. Besides, it can be seen from the figure that PT-SABB incurs smaller communication overheads than AsymPT-FB when solving the problems with low tightness, which demonstrates forward bounding is expensive in message-passing. Additionally, it can be concluded that PT-ISABB algorithms with large k require much more NCLOs than the other competitors when solving problems with low tightness. That is no surprise since inference on problems with large domain sizes would be more expensive, and a search-based algorithm actually can find a feasible solution very quickly even if the lower bounds are poor when solving these problems.

Fig. 8 gives privacy losses under different tightness. Privacy loss in PT-ISABB comes from both the inference phase and the search phase. Specifically, since the variable elimination is performed by parents (i.e., line 4 of the inference phase), parents can easily figure out which pairs of assignments are feasible with respect to the constraints enforced in children sides from the zero entries of the utilities from children. Thus, the inference phase would cause a half privacy loss only on each tree edge in the worst case, which is still better than leaking at least a half privacy if we directly use DPOP to solve problems. Besides, the direct disclosure mechanism of the search phase in which agents request their (pseudo) parents to expose the private costs before expanding the search space also leads to the privacy loss. However, the loss could be much reduced

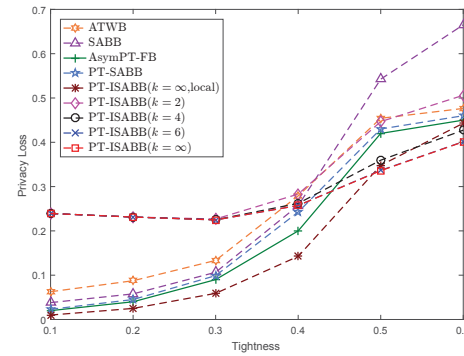


Figure 8: Privacy losses under different tightness

via the effective pruning. It can be seen that our proposed algorithm leaks more privacy than the other competitors when solving the problems with low tightness. That is no surprise because these problems usually have feasible solutions, which leads to the fact that most of entries in a utility from a child are zero. That being said, PT-ISABB with $k \geq 4$ leaks less privacy than the other competitors when solving the problems with high tightness. The reason is twofold: parents can no longer infer the feasible assignment pairs as the tightness grows, and the inference phase produces tight lower bounds which lead to the effective pruning in the search phase. Besides, it is worth mentioning that the local elimination version of PT-ISABB performs better in terms of privacy preservation when solving the problems with low tightness. That is because variables are already eliminated before sending utilities to their parents. As a result, parents can only know the best utilities they can achieve, but cannot figure out the corresponding assignments to their children.

6 CONCLUSION

It is known that DPOP/ADPOP for DCOP cannot be directly applied to ADCOP due to a privacy concern. In this paper, we take ADPOP into solving ADCOP for the first time by combining with a tree-based variation of SABB, and present a two-phase complete algorithm called PT-ISABB. In the inference phase, a non-local elimination version of ADPOP is performed to solve a subset of constraints and build look-up tables for the tighter lower bounds. In the search phase, a tree-based variation of SABB is implemented to exhaust the search space. The experimental results show that our algorithms exhibit great superiorities over state-of-the-art search based algorithms, as well as the local elimination version of PT-ISABB. Also, our algorithms leak less privacy when solving complex problems.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous referees for their valuable comments and helpful suggestions. This work is supported by the Chongqing Research Program of Basic Research and Frontier Technology under Grant No.:cstc2017jcyjAX0030, Fundamental Research Funds for the Central Universities under Grant No.: 2018CDXYJSJ0026, National Natural Science Foundation of China under Grant No.: 51608070 and the Graduate Research and Innovation Foundation of Chongqing, China under Grant No.: CYS18047

REFERENCES

- [1] Ismel Brito, Amnon Meisels, Pedro Meseguer, and Roie Zivan. 2009. Distributed constraint satisfaction with partially known constraints. *Constraints* 14, 2 (2009), 199–234.
- [2] David A Burke, Kenneth N Brown, Mustafa Dogru, and Ben Lowe. 2007. Supply chain coordination through distributed constraint optimization. In *The 9th International Workshop on DCR*.
- [3] Ziyu Chen, Yanchen Deng, and Tengfei Wu. 2017. An Iterative Refined Max-sum_AD Algorithm via Single-side Value Propagation and Local Search. In *Proc. of the 16th Conference on AAMAS*. 195–202.
- [4] Alessandro Farinelli, Alex Rogers, Adrian Petcu, and Nicholas R Jennings. 2008. Decentralised coordination of low-power embedded devices using the max-sum algorithm. In *Proc. of the 7th AAMAS*. 639–646.
- [5] Ferdinando Fioretto, William Yeoh, and Enrico Pontelli. 2016. A dynamic programming-based MCMC framework for solving DCOPs with GPUs. In *International Conference on Principles and Practice of Constraint Programming*. Springer, 813–831.
- [6] Ferdinando Fioretto, William Yeoh, Enrico Pontelli, Ye Ma, and Satishkumar J Ranade. 2017. A Distributed Constraint Optimization (DCOP) Approach to the Economic Dispatch with Demand Response. In *Proc. of the 16th Conference on AAMAS*. 999–1007.
- [7] Amir Gershman, Amnon Meisels, and Roie Zivan. 2009. Asynchronous forward bounding for distributed COPs. *Journal of Artificial Intelligence Research* 34 (2009), 61–88.
- [8] Tal Grinshpoun, Alon Grubshtein, Roie Zivan, Arnon Netzer, and Amnon Meisels. 2013. Asymmetric Distributed Constraint Optimization Problems. *Journal of Artificial Intelligence Research* 47 (2013), 613–647.
- [9] Patricia Gutierrez, Pedro Meseguer, and William Yeoh. 2011. Generalizing ADOPT and BnB-ADOPT. In *Proc. of the 22nd IJCAI*. 554–559.
- [10] Katsutoshi Hirayama and Makoto Yokoo. 1997. Distributed partial constraint satisfaction problem. In *International Conference on Principles and Practice of Constraint Programming*. 222–236.
- [11] Omer Litov and Amnon Meisels. 2017. Forward bounding on pseudo-trees for DCOPs and ADCOPs. *Artificial Intelligence* 252 (2017), 83–99.
- [12] Rajiv T Maheswaran, Jonathan P Pearce, and Milind Tambe. 2004. Distributed Algorithms for DCOP: A Graphical-Game-Based Approach. In *Proceeding of ISCA PDCS'04*. 432–439.
- [13] Rajiv T. Maheswaran, Milind Tambe, Emma Bowring, Jonathan P. Pearce, and Pradeep Varakantham. 2004. Taking DCOP to the Real World: Efficient Complete Solutions for Distributed Multi-Event Scheduling. In *Proc. of the 3rd AAMAS*. 310–317.
- [14] Pragnesh Jay Modi, Wei-Min Shen, Milind Tambe, and Makoto Yokoo. 2005. ADOPT: Asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence* 161, 1-2 (2005), 149–180.
- [15] Tânia L Monteiro, Guy Pujolle, Marcelo E Pellenz, Manoel C Penna, and Richard Demo Souza. 2012. A multi-agent approach to optimal channel assignment in w lans. In *Wireless Communications and Networking Conference (WCNC)*. 2637–2642.
- [16] Arnon Netzer, Alon Grubshtein, and Amnon Meisels. 2012. Concurrent forward bounding for distributed constraint optimization problems. *Artif. Intell.* 193 (2012), 186–216.
- [17] Steven Okamoto, Roie Zivan, and Aviv Nahon. 2016. Distributed breakout: beyond satisfaction. In *Proc. of the 25th IJCAI*. 447–453.
- [18] Brammert Ottens, Christos Dimitrakakis, and Boi Faltings. 2017. DUCT: An Upper Confidence Bound Approach to Distributed Constraint Optimization Problems. *ACM Transactions on Intelligent Systems and Technology (TIST)* 8, 5 (2017), 69.
- [19] Adrian Petcu and Boi Faltings. 2005. Approximations in distributed optimization. In *International Conference on Principles and Practice of Constraint Programming*. 802–806.
- [20] Adrian Petcu and Boi Faltings. 2005. A Scalable Method for Multiagent Constraint Optimization. In *Proc. of the 19th IJCAI*. 266–271.
- [21] Adrian Petcu and Boi Faltings. 2006. ODPOP: an algorithm for open/distributed constraint optimization. In *Proc. of the 21st AAAI*. 703–708.
- [22] Adrian Petcu and Boi Faltings. 2007. MB-DPOP: a new memory-bounded algorithm for distributed optimization. In *Proc. of the 20th IJCAI*. 1452–1457.
- [23] Sarvapali D Ramchurn, Perukrishnen Vytelingum, Alex Rogers, and Nick Jennings. 2011. Agent-based control for decentralised demand side management in the smart grid. In *Proc. of the 10th AAMAS*. 5–12.
- [24] Alex Rogers, Alessandro Farinelli, Ruben Stranders, and Nicholas R Jennings. 2011. Bounded approximate decentralised coordination via the max-sum algorithm. *Artificial Intelligence* 175, 2 (2011), 730–759.
- [25] Evan A Sultanik, Pragnesh Jay Modi, and William C Regli. 2007. On modeling multiagent task scheduling as a distributed constraint optimization problem. In *Proc. of the 20th IJCAI*. 1531–1536.
- [26] Meritxell Vinyals, Juan A Rodriguez-Aguilar, and Jesús Cerquides. 2009. Generalizing DPOP: Action-GDL, a new complete algorithm for DCOPs. In *Proc. of The 8th AAMAS*. 1239–1240.
- [27] William Yeoh, Ariel Felner, and Sven Koenig. 2010. BnB-ADOPT: An asynchronous branch-and-bound DCOP algorithm. *Journal of Artificial Intelligence Research* 38 (2010), 85–133.
- [28] William Yeoh and Makoto Yokoo. 2012. Distributed problem solving. *AI Magazine* 33, 3 (2012), 53.
- [29] Weixiong Zhang, Guandong Wang, Zhao Xing, and Lars Wittenburg. 2005. Distributed stochastic search and distributed breakout: properties, comparison and applications to constraint optimization problems in sensor networks. *Artificial Intelligence* 161, 1-2 (2005), 55–87.
- [30] Roie Zivan and Hilla Peled. 2012. Max/min-sum distributed constraint optimization through value propagation on an alternating DAG. In *Proc. of the 11th AAMAS*. 265–272.