

Privacy Preserving Solution of DCOPs by Local Search

Shmuel Goldklang¹, Tal Grinshpoun², Tamir Tassa¹

¹Department of Mathematics and Computer Science, The Open University of Israel

²Department of Industrial Engineering and Management, Ariel University

shmuel.goldklang@gmail.com, talgr@ariel.ac.il, tamirta@openu.ac.il

Abstract

One of the main reasons for solving constraint optimization problems in a distributed manner is maintaining agents' privacy. Several studies in the past decade devised privacy-preserving versions of Distributed Constraint Optimization Problem (DCOP) algorithms. Some of those algorithms were complete, i.e., finding an optimal solution, while others were incomplete. The main advantage of the incomplete approach is in its scalability to large problems. One of the important incomplete paradigms for solving DCOPs is local search. Yet, so far no privacy-preserving algorithm for solving DCOPs by means of local search was devised. We present P-DSA, a privacy-preserving implementation of the classical local-search algorithm DSA that preserves topology, constraint, and assignment/decision privacy. Comparing its performance to that of P-Max-Sum, which is another privacy-preserving implementation of an incomplete DCOP algorithm, shows that P-DSA is significantly more scalable and issues much better solutions than P-Max-Sum. Therefore, P-DSA emerges as a suitable solution for practitioners addressing large-scale DCOPs with privacy considerations.

1 Introduction

The Distributed Constraint Optimization Problem (DCOP) is a general model for solving distributed combinatorial problems that has a wide range of applications in artificial intelligence. Complete algorithms for DCOP-solving [Modi *et al.*, 2005; Petcu and Faltings, 2005; Gershman *et al.*, 2009] are guaranteed to find the optimal solution, but because DCOPs are NP-hard, these algorithms' worst-case runtime is exponential. Thus, there is a growing interest in incomplete algorithms, which may find sub-optimal solutions but run quickly enough to be applied on large-scale problems or real-time applications [Maheswaran *et al.*, 2004; Zhang *et al.*, 2005; Teacy *et al.*, 2008; Zivan *et al.*, 2014].

Approaches of incomplete DCOP algorithms include inference (Max-Sum [Farinelli *et al.*, 2008]), sampling (DUCT [Ottens *et al.*, 2017], D-Gibbs [Nguyen *et al.*, 2019]), region optimal (KOPT [Katagishi and Pearce,

2007], DALO [Kiekintveld *et al.*, 2010]), and local search (DSA [Zhang *et al.*, 2005], MGM [Maheswaran *et al.*, 2004], and DBA [Hirayama and Yokoo, 2005]). The latter approach is extremely popular due to its simplicity and runtime efficiency.

Privacy is one of the main motivations for solving constraint problems in a distributed manner. Preserving privacy is most important in distributed scenarios in which agents represent people who would not like their personal preferences and actions to be revealed, e.g., meeting scheduling [Gershman *et al.*, 2008], and smart environments (such as smart homes) [Rust *et al.*, 2016; Fioretto *et al.*, 2017]. The term privacy is quite broad, a fact that gave rise to several categorizations of the different types of privacy [Léauté and Faltings, 2013; Greenstadt *et al.*, 2007; Grinshpoun, 2012]. In this paper, we relate to the categorization of Léauté and Faltings [2013] that distinguishes between agent privacy, topology privacy, constraint privacy, and decision privacy.

Most studies that evaluated distributed constraint algorithms in terms of privacy considered complete algorithms [Silaghi and Mitra, 2004; Maheswaran *et al.*, 2006; Greenstadt *et al.*, 2006; Doshi *et al.*, 2008; Léauté and Faltings, 2013; Grinshpoun and Tassa, 2016]. Some work has focused on measuring the extent of constraint privacy loss [Maheswaran *et al.*, 2006; Greenstadt *et al.*, 2006]. Doshi *et al.* [2008] proposed to inject privacy loss as a criterion to the problem-solving process. Some previous work was also directed towards reducing constraint privacy loss. Most efforts in the development of privacy-preserving search algorithms focused on DCSP, which is the *satisfaction* variant of DCOP. Examples include [Nissim and Zivan, 2005; Silaghi and Mitra, 2004; Yokoo *et al.*, 2005]. The work of Silaghi and Mitra [2004] addressed both satisfaction and optimization problems. However, the proposed solution is strictly limited to small-scale problems since it depends on an exhaustive search of all possible assignments. Several privacy-preserving versions of the DPOP algorithm [Petcu and Faltings, 2005] were proposed in the past [Greenstadt *et al.*, 2007; Silaghi *et al.*, 2006], including a more recent study by Léauté and Faltings [2013] that proposed several versions of DPOP that provide strong privacy guarantees. While these versions have been designed for DCSPs, some of them may also be applicable to DCOPs. Explicitly for DCOPs, Grinshpoun and Tassa [2016] and Tassa *et al.* [2021] devised variants of

86 SyncBB [Hirayama and Yokoo, 1997], which preserve topology,
87 constraint, and decision privacy.

88 While the problem sizes for which complete DCOP algorithms
89 are applicable are limited, the problem worsens when privacy-preserving
90 algorithms are considered, due to the substantial runtime overhead that
91 privacy preservation incurs. Consequently, several studies focused on
92 privacy-preserving incomplete algorithms. Tassa et al. [2017] and
93 Kogan et al. [2023] proposed variations of an incomplete inference
94 algorithm, Max-Sum [Farinelli et al., 2008], which preserve topology,
95 constraint, and decision privacy. Additionally, Grinshpoun et al. [2019]
96 devised an incomplete region-optimal algorithm that preserves constraint
97 privacy and partial decision privacy. However, though incomplete, the above
98 algorithms are very elaborate and are inapplicable to large-scale
99 problems. Specifically, the runtime of the Max-Sum-based algorithms is
100 exponential in the *arity* of the constraints, which makes them unsuitable
101 for problems with global constraints, e.g., satellite scheduling [Krigman
102 et al., 2024] and course allocation [Khakhiashvili et al., 2021].

103 Recently, Vion et al. [2022] proposed a local search algorithm that
104 *controls* the loss of domain privacy [Grinshpoun, 2012] by following the
105 Utilitarian DCOP model [Doshi et al., 2008; Savaux et al., 2020], in
106 which privacy loss is traded off with solution quality. However, their
107 approach is only relevant in the Open Constraints Programming model
108 [Faltings and Macho-Gonzalez, 2005], where the domains are not known
109 in advance and grow as the solving process advances.

110 **Our contributions.** We present here P-DSA, a privacy-preserving
111 implementation of the classical local-search algorithm DSA. We show that
112 it offers topology privacy, constraint privacy, and assignment/decision
113 privacy. We compare its performance to that of P-Max-Sum [Tassa et al.,
114 2017], a privacy-preserving implementation of the incomplete DCOP
115 algorithm Max-Sum [Farinelli et al., 2008] which also protects topology,
116 constraint and assignment/decision information. We show that P-DSA is
117 significantly more scalable and issues much better solutions than P-Max-
118 Sum. In fact, while P-DSA was able to solve in short time (3 minutes)
119 problems involving as high as 100 agents, prior studies on privacy-
120 preserving DCOP algorithms report experiments with at most 24 agents
121 and runtimes that are significantly higher.¹ Therefore, P-DSA emerges
122 as a suitable choice for solving large-scale DCOPs in a privacy-
123 preserving manner.

130 2 DCOP background

A Distributed Constraint Optimization Problem (DCOP, [Hirayama and Yokoo, 1997]) is a tuple $\langle \mathcal{A}, \mathcal{X}, \mathcal{D}, \mathcal{R} \rangle$ where $\mathcal{A} = \{A_1, \dots, A_n\}$ is a set of agents, $\mathcal{X} = \{X_1, \dots, X_n\}$ is a set of variables, $\mathcal{D} = \{D_1, \dots, D_n\}$ is a set of finite domains, and \mathcal{R} is a set of relations (constraints). Each variable X_i takes values in the domain D_i , and it is held by the agent

¹To the best of our knowledge, the only exception is the work of Damle et al. [2024] that presented P-Gibbs, which is a differentially private implementation of SD-Gibbs [Nguyen et al., 2019]. However, differential privacy is a paradigm that is based on injecting random noise; hence it is not directly comparable to the cryptographic paradigm that does not alter the outputs of the underlying algorithm.

A_i . Each constraint $C \in \mathcal{R}$ defines for a given pair of variables some non-negative cost; formally, a constraint takes the form $C_{i,j} : D_i \times D_j \rightarrow [0, q]$, for some $1 \leq i \leq j \leq n$, where q is a publicly known maximal constraint cost. (Note that if $i = j$ then the constraint is unary.) The goal in constraint optimization problems is to find an assignment of values to all n variables,

$$(X_1, \dots, X_n) \leftarrow \mathbf{x} := (x_1, \dots, x_n) \in \mathbf{D} := D_1 \times \dots \times D_n,$$

such that the overall incurred cost $\sum_{C_{i,j} \in \mathcal{R}} C_{i,j}(x_i, x_j)$ is minimal.

Our framework can also include the case of *hard* constraints, i.e., combinations of assignments that are strictly forbidden, see [Kumar et al., 2008]. Our framework is the one that is studied in most prior art. Some studies consider extensions to this framework by (a) assuming that each agent may hold more than one variable [Yokoo and Hirayama, 2000; Burke and Brown, 2006; Grinshpoun, 2015; Fioretto et al., 2016], (b) including constraints of arity greater than two [Kim and Lesser, 2013], and (c) assuming asymmetric constraints that incur different costs to each of the involved agents [Grinshpoun et al., 2013]. However, here we focus on the framework as defined above, which already introduces the main challenges of DCOPs.

Léauté and Faltings [2013] distinguished between four notions of privacy: agent privacy (who are the agents in the problem setting), topology privacy (hiding information on the constraint graph), constraint privacy (hiding information on the costs in the constraints), and assignment/decision privacy (protecting the intermediate/final assignments).

2.1 The Distributed Stochastic Algorithm

Here we describe the classic local search DCOP algorithm that was presented by Zhang et al. [2005] – the Distributed Stochastic Algorithm (DSA). We start by introducing a key notion in local search algorithms:

Definition 1 (Neighborhood). *The neighborhood of agent A_i is the set of all agents that are constrained with A_i , i.e., $N(A_i) := \{A_j \in \mathcal{A} : \exists C_{i,j} \in \mathcal{R}\}$. The complete neighborhood of A_i is $N^+(A_i) := N(A_i) \cup \{A_i\}$.*

Algorithm 1: The DSA algorithm

```

1 forall  $i \in [n]$  do
2    $A_i$  selects at random  $x_i \in D_i$ 
3 forall  $\ell = 1, \dots, L$  do
4   forall  $i \in [n]$  do
5      $A_i$  sends  $x_i$  to all  $A_j \in N(A_i)$ 
6     forall  $i \in [n]$  do
7        $A_i$  samples uniformly at random a real
8          $x \in [0, 1]$ 
9       if  $x \leq p$  then
10         $A_i$  chooses  $y_i \in D_i$  that minimizes
            $\sum_{A_j \in N(A_i)} C_{i,j}(y_i, x_j)$ 
            $A_i$  updates  $x_i \leftarrow y_i$ 

```

Algorithm 1 describes DSA. The algorithm starts by generating an initial random assignment $\mathbf{a} \in \mathbf{D}$ (Lines 1-2).² It then keeps updating that assignment by performing a preset number of iterations L (Lines 3-10). The assignment in the final iteration is the algorithm's output. Each iteration starts with each agent updating its neighbors on its current assignment (Lines 4-5). Then, each agent is allowed, with probability p , to change its local assignment to the best possible value (Lines 6-10).

The utilization of the stochastic factor p enables DSA to escape local minima and avoid infinite loops. However, it renders DSA non-monotone in the sense that the cost of the solution in one iteration is not necessarily smaller than the cost of the solution in the previous iteration. It is possible to enhance DSA with a so-called *anytime mechanism* [Zivan *et al.*, 2014]. Such a mechanism finds the best solution visited throughout the run of the algorithm. In general, in order to report the best solution visited, the algorithm needs to compute the overall cost after each iteration, and if that overall cost is the minimum so far, record that cost and the corresponding assignment.

3 Cryptographic background

Here, we briefly describe the cryptographic machinery we use in our protocols. In Section 3.1 we discuss threshold secret sharing, and then, in Section 3.2, we describe secure computations over secret-shared values.

3.1 Shamir's secret sharing

Secret sharing schemes [Shamir, 1979] are protocols that enable distributing a secret scalar s among a set of agents, A_1, \dots, A_n . Each agent, A_h , $h \in [n]$, gets a random value $[[s]]_h$, called a *share*, so that some subsets of those shares enable the reconstruction of s , while each of the other subsets of shares reveals no information on s . In its most basic form, called *Threshold Secret Sharing*, there is a threshold value $t \leq n$, and then a subset of shares enables the reconstruction of s iff its size is at least t .

Shamir's t -out-of- n threshold secret sharing scheme [Shamir, 1979] operates over a finite field \mathbf{Z}_q , where $q > n$ is a prime sufficiently large so that all possible secrets may be represented in \mathbf{Z}_q . It has two procedures: **Share** and **Reconstruct**:

- **Share** $_{t,n}(s)$. The procedure samples a uniformly random polynomial $f(\cdot)$ over \mathbf{Z}_q , of degree at most $t - 1$, where the free coefficient is the secret s . That is, $f(x) = s + a_1x + a_2x^2 + \dots + a_{t-1}x^{t-1}$, where a_j , $1 \leq j \leq t - 1$, are selected independently and uniformly at random from \mathbf{Z}_q . The procedure outputs n values, $[[s]]_h = f(h)$, $h \in [n]$, where $[[s]]_h$ is the share given to A_h . The entire set of shares, denoted $[[s]] := \{[[s]]_h : h \in [n]\}$, is called a (t, n) -sharing of s .

- **Reconstruct** $_t([[s]])$. The procedure is given any selection of t shares out of the (t, n) -sharing of s . It then interpolates a polynomial $f(\cdot)$ of degree at most $t - 1$ using the given points and outputs $s = f(0)$. Any selection of t shares will yield the secret s , as t points determine a unique polynomial of degree

at most $t - 1$. On the other hand, any selection of $t - 1$ shares or less reveals nothing about the secret s .

Hereinafter, we set the secret sharing threshold to be

$$t := \lfloor (n + 1)/2 \rfloor. \quad (1)$$

Namely, to reconstruct the secret, at least half of the agents must collaborate. Hence, if the set of n agents has an honest majority (in the sense that more than $n/2$ agents would not try to combine their shares in order to recover secret values), all shared values will remain fully protected.

In what follows, we shall use the following terminology and notations. Let s be a secret known to some agent A_i , $i \in [n]$. Then if A_i performs the procedure $\text{Share}_{t,n}(s)$, we will simply say that A_i distributes a (t, n) -sharing of s .

If the agents have a (t, n) -sharing $[[s]]$ in some secret s and they wish to let one of them, say A_i , reconstruct the secret s , then at least $t - 1$ agents would send their shares to A_i who will proceed to apply the **Reconstruct** procedure on the t shares it has. We will describe this procedure shortly by writing $s \leftarrow \text{Reconstruct}([[s]]; A_i)$.

If $\mathbf{s} = (s_1, \dots, s_m) \in \mathbf{Z}_q^m$ is a vector of secrets held by A_i , then by saying that A_i distributes a (t, n) -sharing of \mathbf{s} we mean that A_i distributes a (t, n) -sharing of each of the entries of \mathbf{s} , independently.

Let $a \in \mathbf{Z}_q$ be any value that is publicly known to all agents. Then by $[[a]]$ we mean the set of (t, n) -shares $\{[[a]]_h = a : h \in [n]\}$. It is easy to see that this set of shares indeed defines a unique polynomial of degree at most $t - 1$, which is the constant polynomial $f(\cdot) \equiv a$, and therefore it is a proper (t, n) -sharing of the value a . Such a sharing does not require any communication between the agents nor any polynomial computations, since a is publicly known.

Let $a, b, c \in \mathbf{Z}_q$ be three values publicly known to all, and let s and s' be two secrets in which the agents already hold (t, n) -sharings, denoted $[[s]]$ and $[[s']]$. Then

$$a + b[[s]] + c[[s']] := \{a + b[[s]]_h + c[[s']]_h : h \in [n]\} \quad (2)$$

is a proper (t, n) -sharing of $\hat{s} := a + bs + cs'$, and its computation needs no interaction between the agents, thanks to the affinity of secret sharing. By writing

$$[[\hat{s}]] \leftarrow a + b[[s]] + c[[s']]$$

we mean that each agent A_h , $h \in [n]$, sets $[[\hat{s}]]_h \leftarrow a + b[[s]]_h + c[[s']]_h$, so that now the agents hold a (t, n) -sharing of $\hat{s} = a + bs + cs'$ without needing to interact or perform any further polynomial computations.

3.2 Secure computations over secret sharings

Let a and b be two secret values in the field \mathbf{Z}_q , and assume that A_1, \dots, A_n hold (t, n) -sharings in them, denoted $[[a]] = \{[[a]]_h : h \in [n]\}$ and $[[b]] = \{[[b]]_h : h \in [n]\}$. A secure multiplication protocol is a protocol of the form

$$[[c]] \leftarrow \text{SecureMult}([[a]], [[b]]), \quad (3)$$

that takes the (t, n) -sharings of a and b and computes from them a (t, n) -sharing of $c = a \cdot b$ in a secure manner, namely, without revealing to the agents any information on a , b , or $c = ab$. Damgård and Nielsen [2007] designed such a secure

²Throughout this paper, for any integer n , $[n] := \{1, \dots, n\}$.

261 multiplication protocol. In our experiments, we used that pro-
 262 tocol with the performance improvements that were proposed
 263 by Chida et al. [2018].

264 Another computation on secret shares that we will need is
 265 secure comparison. Under the same assumptions as above, a
 266 secure comparison protocol is a protocol of the form

$$[[c]] \leftarrow \text{SecureCompare}([[a]], [[b]]), \quad (4)$$

267 that takes the (t, n) -sharings of a and b and computes from
 268 them a (t, n) -sharing of $c = 1_{a < b}$, where hereinafter if \mathcal{P} is
 269 a predicate then $1_{\mathcal{P}}$ is a bit that equals 1 if the predicate \mathcal{P}
 270 holds and equals 0 otherwise. As before, such a protocol is
 271 secure in the sense that it does not reveal to the agents any
 272 information on a , b , or $c = 1_{a < b}$. Nishide and Ohta [2007]
 273 proposed such a secure comparison protocol.

274 4 Private DSA

275 In this section, we describe Private DSA (P-DSA), an im-
 276 plementation of DSA that preserves topology, constraint, and
 277 decision privacy. In order to achieve those privacy goals, P-
 278 DSA employs the following principles:

279 (1) To achieve topology privacy, every pair of agents that
 280 are not constrained creates a zero constraint matrix between
 281 themselves, and, subsequently, the algorithm acts on a com-
 282 plete constraint graph. None of the other agents is able to dis-
 283 tinguish between fake constraint matrices (i.e., zero matrices)
 284 and genuine ones due to the next principle in P-DSA's design,
 285 which distinguishes its operation from that of the basic DSA.

286 (2) To achieve constraint privacy, all constraint matrices
 287 are secret-shared among all agents, and all computations that
 288 rely on those matrices use the shares rather than the actual
 289 constraint matrices.

290 (3) To achieve decision privacy, in each iteration of the al-
 291 gorithm whenever an agent selects an assignment to its vari-
 292 able, it does not send that assignment to its neighbors; instead,
 293 it secret shares information on the costs that such an assign-
 294 ment incurs vis-a-vis each of the other agents.

295 The latter principle raises a considerable computational
 296 challenge: how can each of the agents perform the compu-
 297 tations that DSA mandates when it does not know the current
 298 assignments of its neighboring agents? We tackle that chal-
 299 lenge by designing multi-party sub-protocols to be run jointly
 300 by all agents. In those collaborative sub-protocols, all agents
 301 use the secret shares they hold in order to enable each agent
 302 to compute the next assignment from its domain. In doing so,
 303 none of the agents get any wiser about that assignment or any
 304 other private information.

305 We assume hereinafter that all agents know the sizes of all
 306 domains, namely, $m_i := |D_i|$ for all $i \in [n]$. Moreover, each
 307 agent A_i , $i \in [n]$, generates an ordering of the values in its
 308 domain, $D_i = \{a_1^i, \dots, a_{m_i}^i\}$, and publishes that ordering to
 309 each of its neighbors, $A_j \in N(A_i)$. Therefore, each con-
 310 straint $C_{i,j}$ can be described as a matrix of m_i rows and m_j
 311 columns, where $C_{i,j}(r, s)$ equals the value of the constraint
 312 when $X_i = a_r^i$ and $X_j = a_s^j$. In what follows, we will think
 313 of $C_{i,j}$ as a matrix rather than a function over $D_i \times D_j$.

314 Protocol 2 describes P-DSA — a private implementation of
 315 DSA. First, each agent A_i selects a random assignment to its

variable. A_i does that by selecting a random index $r_i \in [m_i]$,
 316 and then the corresponding assignment to X_i is $a_{r_i}^i$, $i \in [n]$
 317 (Lines 1-2).
 318

The main loop takes place in Lines 3-15. First, each agent
 319 A_i , $i \in [n]$, secretly shares its current assignment, $a_{r_i}^i$, with
 320 all agents. To do that, A_i distributes to all agents (t, n) -shares
 321 in the r_i -th row in each of the constraint matrices that it has
 322 vis-a-vis each of the other $n - 1$ agents (namely, also with
 323 agents outside its neighborhood). Let $\mathbf{w}_{i,j}$ denote the r_i -th
 324 row in the constraint matrix $C_{i,j}$, for some $j \in [n] \setminus \{i\}$, i.e.,
 325

$$\mathbf{w}_{i,j} = (\mathbf{w}_{i,j}(u) : u \in [m_j]),$$

$$\text{where } \mathbf{w}_{i,j}(u) = C_{i,j}(r_i, u), u \in [m_j]. \quad (5)$$

A_i distributes (t, n) -shares in each of the m_j entries of that
 326 vector, where the sharing of $\mathbf{w}_{i,j}(u)$ is denoted $[[\mathbf{w}_{i,j}(u)]] =$
 327 $\{[\mathbf{w}_{i,j}(u)]_h : h \in [n]\}$, while the sharing of the entire vec-
 328 tor is denoted $[[\mathbf{w}_{i,j}]]$. The overall number of scalars that A_i
 329 shares at this stage (Lines 4-6) is $\sum_{j \in [n] \setminus \{i\}} m_j$.
 330

We would like to clarify that the secret sharing done in
 331 Lines 4-6 is excessive. Indeed, if $a \neq b \in [n]$ then the scalar
 332 $C_{a,b}(r_a, r_b)$ is shared when $i = a$ and $j = b$, as it is in the r_a -
 333 th row of the matrix $C_{a,b}$, but also when $i = b$ and $j = a$, as
 334 it is in the r_b -th row of the matrix $C_{b,a}$ which is the transpose
 335 of $C_{a,b}$. However, this excessive secret sharing will pay off
 336 later on in the computation.
 337

Before moving on, let us fix $i \in [n]$ and $j \in [n] \setminus \{i\}$.
 338 Then for any $u \in [m_i]$, $\mathbf{w}_{j,i}(u)$ is the cost that A_i would pay
 339 if it sets $X_i = a_u^i$, given the current assignment of A_j to its
 340 variable, $X_j = a_{r_j}^j$. Therefore, if we define
 341

$$\mathbf{w}_i(u) := \sum_{j \in [n] \setminus \{i\}} \mathbf{w}_{j,i}(u), \quad u \in [m_i], \quad (6)$$

we have by Eq. (5) and the symmetry of the constraints (in
 342 the sense that $C_{i,j} = C_{j,i}^T$),
 343

$$\mathbf{w}_i(u) = \sum_{j \in [n] \setminus \{i\}} C_{j,i}(r_j, u) = \sum_{j \in [n] \setminus \{i\}} C_{i,j}(u, r_j), \quad u \in [m_i]. \quad (7)$$

Hence, $\mathbf{w}_i(u)$ is the overall cost for A_i if it sets $X_i = a_u^i$,
 344 given the current assignments that all other agents have for
 345 their variables. In Lines 7-9 all agents compute (t, n) -shares
 346 in $\mathbf{w}_i(u)$ for all $i \in [n]$ and for all $u \in [m_i]$. Note that
 347 it is a local computation that does not require the agents to
 348 communicate.
 349

Next, the main task of each agent A_i is to find the best
 350 assignment to its variable given the current assignments of
 351 all neighboring variables (as encoded in the secret shares that
 352 all agents have distributed in Lines 4-6) and storing the in-
 353 dex of that assignment in r_i . However, we recall that such
 354 a computation takes place only in probability p , while oth-
 355 erwise, in probability $1 - p$, A_i retains its current assign-
 356 ment. Hence, A_i starts by generating a uniformly random
 357 real number $x \in [0, 1]$ (Line 11), and only if $x \leq p$ it pro-
 358 ceeds to the computational task of finding the best assignment
 359 for its variable, given the current assignments of its neigh-
 360 boring agents. That computation is carried out in the sub-
 361 protocol FindBestAssignment (Line 13). In that sub-protocol,
 362 the agents jointly and securely compute a (t, n) -sharing of the
 363

364 index $k_i \in [m_i]$ of the currently best assignment to X_i from
 365 D_i . After its completion, all agents send to A_i their shares in
 366 k_i , and A_i proceeds to recover k_i (Line 14) and store it in r_i
 367 (Line 15).

368 After performing L such iterations (Lines 3-15), each of
 369 the agents stores the last assignment to its variable (Lines 16-
 370 17).

Protocol 2: P-DSA – Private DSA

```

1 forall  $i \in [n]$  do
2    $A_i$  selects at random  $r_i \in [m_i]$ 
3 forall  $\ell = 1, \dots, L$  do
4   forall  $i \in [n]$  do
5     forall  $j \in [n] \setminus \{i\}$  do
6        $A_i$  distributes a  $(t, n)$ -sharing of  $[[\mathbf{w}_{i,j}]]$ 
7     forall  $i \in [n]$  do
8       forall  $u \in [m_i]$  do
9          $[[\mathbf{w}_i(u)]] \leftarrow \sum_{j \in [n] \setminus \{i\}} [[\mathbf{w}_{j,i}(u)]]$ 
10    forall  $i \in [n]$  do
11       $A_i$  samples uniformly at random  $x \in [0, 1]$ 
12      if  $x \leq p$  then
13        FindBestAssignment( $i; [[k_i]]$ )
14         $k_i \leftarrow \text{Reconstruct}([[k_i]]; A_i)$ 
15         $A_i$  sets  $r_i \leftarrow k_i$ 
16 forall  $i \in [n]$  do
17    $A_i$  sets  $X_i \leftarrow a_{r_i}^i$ 

```

371 **4.1 The sub-protocol FindBestAssignment**

372 Here, we describe Sub-protocol 3, called FindBestAssign-
 373 ment. The sub-protocol, which is executed by all agents,
 374 scans the values in X_i 's domain, $D_i = \{a_u^i : u \in [m_i]\}$,
 375 and computes a (t, n) -sharing $[[k_i]]$ in the index $k_i \in [m_i]$
 376 that issues the currently minimal aggregated cost for A_i .

377 Before describing the computations in the sub-protocol, we
 378 make the following observations. Let c_i and c_j be two in-
 379 dexed scalars, where $i < j$. Then

$$\min(c_i, c_j) = c_i + 1_{c_j < c_i} \cdot (c_j - c_i) \quad (8)$$

380 and

$$\arg \min(c_i, c_j) = i + 1_{c_j < c_i} \cdot (j - i) \quad (9)$$

(by $\arg \min$ we mean the smallest index in which the min-
 imum is attained). Hence, if the agents hold (t, n) -shares
 in c_i and in c_j , they can jointly compute (t, n) -shares in
 $\min(c_i, c_j)$ and in $\arg \min(c_i, c_j)$, without learning any in-
 formation on c_i and c_j , by invoking the secure comparison
 and multiplication protocols from Section 3.2. Specifically,
 they will first run

$$[[\beta]] \leftarrow \text{SecureCompare}([[c_j]], [[c_i]])$$

(see Eq. (4)) so that they will hold (t, n) -shares in the bit
 $\beta := 1_{c_j < c_i}$. Then they will run the secure multiplication
 protocol (see Eq. (3)),

$$[[\gamma]] \leftarrow \text{SecureMult}([[c_j]] - [[c_i]], [[\beta]])$$

to get (t, n) -shares in $\gamma := 1_{c_j < c_i} \cdot (c_j - c_i)$. Finally, each
 agent A_h , $h \in [n]$, will compute

$$[[w]]_h \leftarrow [[c_i]]_h + [[\gamma]]_h.$$

In view of Eq. (8), the set $[[w]] = \{[[w]]_h : h \in [n]\}$ is a
 (t, n) -sharing of $w := \min(c_i, c_j)$. In the process of com-
 puting those shares, the agents remain completely oblivious
 to the values of c_i , c_j , and w . A similar course of action can
 issue to the agents a (t, n) -sharing of $\arg \min(c_i, c_j)$, using
 Eq. (9).

We now turn to Sub-protocol 3. Its input is the index i
 of the agent who looks for the currently best assignment to
 its variable. Recall that FindBestAssignment is invoked from
 Protocol 2 in Line 13. At that stage in Protocol 2, all agents
 hold (t, n) -shares in $\mathbf{w}_i(u)$ for all $i \in [n]$ and all $u \in [m_i]$,
 being the aggregated cost for A_i if it sets $X_i \leftarrow a_u^i$, given the
 current assignments to the variables held by its neighbors.

The sub-protocol scans A_i 's domain, D_i , and updates two
 values: k_i that will hold the index of the currently best assign-
 ment and w_i that will hold the corresponding cost. Those two
 values will not be computed explicitly; instead, the agents
 will hold secret shares in them.

Initially (Lines 1-2), the agents set $k_i = 1$ and $w_i = \mathbf{w}_i(1)$.
 Since the agents already hold a secret sharing of the latter
 value, they simply set $[[w_i]]_h = [[\mathbf{w}_i(1)]]_h$, $h \in [n]$. As for
 $k_i = 1$, since it is a publicly known value, then, in view of our
 discussion in Section 3.1, each agent sets $[[k_i]]_h = 1$, $h \in [n]$.

Next, the agents scan the remaining values in D_i (Lines
 3-8). First, they compute shares in $\beta := 1_{\mathbf{w}_i(u) < w_i}$, using
 SecureCompare (see Eq. (4)), in order to compare w_i , the
 minimum found so far, to the cost of the next assignment,
 $\mathbf{w}_i(u)$ (Line 4). Then, they use SecureMult (see Eq. (3)) to
 compute shares in $\gamma := \beta \cdot (\mathbf{w}_i(u) - w_i)$ and in $\delta := \beta \cdot (u - k_i)$
 (Lines 5-6). (Recall that since u is a publicly known value,
 each agent A_h , $h \in [n]$, sets locally $[[u]]_h = u$.) Finally,
 they update the shares in w_i and k_i using Eqs. (8) and (9),
 respectively (Lines 7-8). At the end of the loop, k_i equals the
 index of the best assignment, and w_i equals the associated
 cost. Since P-DSA needs only $[[k_i]]$, the sub-protocol issues
 that sharing as its output.

Comment. The computation of w_i (Lines 5+7) is needed
 for the computation of β (Line 4) in the subsequent iteration,
 a value that is used in updating k_i (Lines 6+8). Hence, since
 w_i is not a desired output of the sub-protocol, it is possible to
 skip Lines 5+7 in the last iteration ($u = m_i$).

Sub-protocol 3: FindBestAssignment – Computing a
 (t, n) -sharing of the index k_i of the currently best as-
 signment for X_i .

Input: i – the index of agent A_i

```

1 forall  $h \in [n]$  do
2    $A_h$  sets  $[[k_i]]_h \leftarrow 1$  and  $[[w_i]]_h \leftarrow [[\mathbf{w}_i(1)]]_h$ 
3 forall  $u = 2, \dots, m_i$  do
4    $[[\beta]] \leftarrow \text{SecureCompare}([[w_i(u)]], [[w_i]])$ 
5    $[[\gamma]] \leftarrow \text{SecureMult}([[c_j]], [[\mathbf{w}_i(u)]] - [[w_i]])$ 
6    $[[\delta]] \leftarrow \text{SecureMult}([[c_j]], [[u]] - [[k_i]])$ 
7    $[[w_i]] \leftarrow [[w_i]] + [[\gamma]]$ 
8    $[[k_i]] \leftarrow [[k_i]] + [[\delta]]$ 

```

Output: A (t, n) -sharing of $[[k_i]]$

4.2 Privacy

Protocol 2 preserves topology, constraint, and assignment/decision privacy, owing to the cryptographic machinery that we use – see Theorem 1. It does not respect agent privacy since it requires all n agents to have a full communication network between them.

Theorem 1. *Under the assumption of honest majority, Protocol 2 preserves topology, constraint, and assignment/decision privacy.*

Proof. The honest majority assumption means that if there exist agents that will try combining their shares in attempt to recover some of the secret-shared values, their number will be smaller than the threshold $t = \lfloor (n + 1)/2 \rfloor$, see Eq. (1). Shamir’s secret sharing scheme is perfect, in the sense that any number of shares smaller than the threshold exposes zero information on the shared secret [Shamir, 1979]. Therefore, the secret shares in each of the private values that are secret-shared during P-DSA reveal no information on the underlying private value. Apart from secret sharing, the agents engage also in multi-party protocols for performing secure multiplication and secure comparison, see Eqs. (3) and (4). The protocols that we use are information-theoretic secure, see [Damgård and Nielsen, 2007; Chida *et al.*, 2018; Nishide and Ohta, 2007]. Given all of the above, it follows the P-DSA fully preserves all constraint information under the honest majority assumption; hence, it offers constraint privacy.

P-DSA operates over a complete constraint graph, in which every pair of agents has a constraint matrix between them. Since all matrices are secret-shared using the threshold t in Eq. (1), which guarantees perfect privacy under the assumption of honest majority, zero matrices are indistinguishable from matrices that represent actual constraints. Therefore, P-DSA offers also topology privacy.

As also all indices of all assignments are encoded through secret shares, we infer that all assignment information, as well as the final decisions, remain fully protected. Hence, P-DSA offers also assignment/decision privacy.

Note that while Protocol 2 hides from each agent the sequence of assignments of other agents, it does reveal to each agent its own sequence of assignments. Protocol 2 can be further enhanced to also hide from each agent the sequence of value assignments to its own variable, including the initial random value assignment. Due to space limitations, we omit the details of this enhancement.

5 Experiments

We implemented P-DSA and compared its performance to P-Max-Sum [Tassa *et al.*, 2017], which is a privacy-preserving implementation of an incomplete DCOP algorithm (Max-Sum [Farinelli *et al.*, 2008]).

Experiments were conducted on a machine equipped with an Intel i5-10400 CPU @ 2.90GHz, 2904 Mhz, 6 Core(s), 12 Logical Processor(s), 16GB DDR4 RAM. The system ran Microsoft Windows 10 Pro, and the code was written in Java 1.8.0 using the SinAlgo sim-

ulation framework. The source code is available on <https://github.com/dcop2025/dcop-sim/tree/main>.

P-DSA was implemented over \mathbf{Z}_q with $q = 2^{31} - 1$. P-Max-Sum was implemented with 512-bit homomorphic encryption.

In our experiments, we compared the quality of the solutions issued by each of those two algorithms within a given time frame. We used the following settings of the main parameters that affect the algorithms’ runtimes:

- Number of agents $n \in \{10, 20, \underline{30}, 40, \dots, 100\}$.
- Domains’ size $m \in \{5, \underline{10}, 15, 20, 25\}$. For simplicity, we assumed that all domains have the same size m .
- Constraint density, $d \in \{0.2, \underline{0.4}, 0.6, 0.8, 1.0\}$ — the fraction of constrained pairs of variables out of all $\binom{n}{2}$ pairs.

To test the effect of each of those three parameters, we set the other two to the value that is underlined in their respective set of tested values and varied the value of the tested parameter. For example, in testing the effect of the number of agents, we set all domain sizes to be $m = 10$ and used constraint density of $d = 0.4$ and then ran experiments with $n \in \{10, \dots, 100\}$.

We refer to each triple $\langle n, m, d \rangle$ as a *configuration*. In each tested configuration, we evaluated both algorithms in the following manner: We selected a new random problem (where a problem consists of the constraint graph as well as the constraint matrices), ran both algorithms on the same problem, and evaluated the cost of their output after $T = 1, 2, 3$ minutes of execution. We repeated that experiment 20 times, and we report the average of the costs obtained by each of the two algorithms within each of the prescribed time frames.

In one set of experiments we used random constraint graphs, where each graph is a random graph of n nodes in which each pair of nodes is connected by an edge in probability d . In another set of experiments we generated scale-free random graphs [Barabási and Albert, 1999] with an initial clique of size 5, and 4 backward edges for each additional node. In all experiments, each constraint matrix was a random $m \times m$ matrix with entries that distribute uniformly on the interval $[0, 10]$.

Number of agents in random graphs. We compared the average cost of solutions issued by each of the two algorithms within each of the three prescribed time frames for a varying number of agents n (where in all problems, the domain size was $m = 10$ and the network density was $d = 0.4$). Table 1 shows the average costs issued by the two algorithms (rounded to the nearest integer). The symbol \perp indicates that the algorithm did not manage to complete even one iteration within the time frame.

We see the overwhelming advantages of P-DSA over P-Max-Sum in terms of scalability and quality of solutions. Indeed, while P-Max-Sum could not produce a solution within 1 minute already for $n = 40$ and could not produce a solution within 3 minutes for $n \geq 60$, P-DSA was able to produce solutions within 1 minute for all $n \leq 60$ and managed to produce a solution within 3 minutes for all tested values of n . Furthermore, the solutions produced by P-DSA were better than those issued by P-Max-Sum by more than 50%

T	$n = 10$	20	30	40	50	60	70	80	90	100
1	27 66	207 330	591 783	1280 ⊥	2197 ⊥	3127 ⊥	⊥ ⊥	⊥ ⊥	⊥ ⊥	⊥ ⊥
2	27 63	186 333	526 733	1121 1433	1944 ⊥	3023 ⊥	4361 ⊥	5796 ⊥	⊥ ⊥	⊥ ⊥
3	27 59	178 318	494 763	1058 1362	1822 2299	2902 ⊥	4112 ⊥	5752 ⊥	7354 ⊥	9087 ⊥

Table 1: Average costs obtained by P-DSA (left in each table cell) and P-Max-Sum (right) for problems in random graphs over a varying number n of agents, within time frames of $T = 1, 2, 3$ minutes. The symbol \perp indicates that the algorithm did not manage to complete even a single iteration within the time frame.

T	$n = 10$	20	30	40	50	60	70	80	90	100
1	55 101	503 661	1496 ⊥	3083 ⊥	5211 ⊥	7732 ⊥	⊥ ⊥	⊥ ⊥	⊥ ⊥	⊥ ⊥
2	53 118	467 683	1386 1671	2836 ⊥	4897 ⊥	7482 ⊥	10717 ⊥	14100 ⊥	⊥ ⊥	⊥ ⊥
3	53 106	463 677	1325 1671	2740 ⊥	4652 ⊥	7260 ⊥	10321 ⊥	14069 ⊥	18095 ⊥	⊥ ⊥

Table 2: Similar to Table 1 but with scale-free graphs.

535 for $n = 10$ and by 20% for the largest problem in which P-
536 Max-Sum issued a solution. In addition, we see that P-DSA
537 always improves the quality of its output when allowed to run
538 for more time, while P-Max-Sum sometimes fluctuates (see,
539 e.g., its outputs when $n = 30$). That is why it is sometimes
540 executed with the anytime mechanism [Zivan *et al.*, 2014]
541 that outputs the best solution visited throughout the run of
542 the algorithm. Such a mechanism has its overhead, and in
543 P-DSA, it appears that there is less need to apply it. (It is im-
544 portant to stress that P-DSA and P-Max-Sum issue the very
545 same intermediate and final assignments as DSA and Max-
546 Sum, respectively. Namely, the cryptographic layer protects
547 the underlying private information but does not alter it.)

T	$m = 5$	10	15	20	25
1	560 714	591 783	669 ⊥	662 ⊥	715 ⊥
2	550 722	526 733	570 830	587 ⊥	619 ⊥
3	550 689	494 763	516 830	527 834	550 ⊥

Table 3: Average costs obtained by P-DSA (left in each table cell) and P-Max-Sum (right) for problems in random graphs over a varying domain size m , within time frames of $T = 1, 2, 3$ minutes.

T	$d = 0.2$	0.4	0.6	0.8	1.0
1	254 316	591 783	980 ⊥	1353 ⊥	1784 ⊥
2	198 317	526 733	889 1212	1244 1641	1645 2053
3	174 327	494 763	843 1179	1197 1588	1575 2023

Table 4: Average costs obtained by P-DSA (left in each table cell) and P-Max-Sum (right) for problems in random graphs over a varying constraint density d , within time frames of $T = 1, 2, 3$ minutes.

558 issued outputs, those of P-DSA had costs lower than those of
559 P-Max-Sum, with improvements ranging from 22% to 38%.

Constraint density in random graphs. Here we fixed $n =$
560 30 and $m = 10$ and varied the constraint density d . The
561 results are given in Table 4. As before, P-DSA issues so-
562 lutions with costs that are significantly lower than P-Max-
563 Sum’s (where in one configuration, the improvement was as
564 high as 47%). As for scalability, P-DSA’s runtime does not
565 depend on the network density since it operates on the com-
566 plete graph, where non-constrained pairs of agents are con-
567 nected by an edge with a zero constraint matrix. P-Max-Sum,
568 on the other hand, works on the original constraint graph,
569 and therefore, its runtime does depend on the network den-
570 sity. Hence, it failed to issue an output on dense networks for
571 which P-DSA did issue an output.

572 Due to lack of space we omit description of experiments
573 that compare the runtimes of P-DSA and the basic DSA, in
574 order to illustrate the price of privacy. We intend to include
575 such experiments in the full version of this study.
576

6 Conclusion

577 We presented here P-DSA – the first privacy-preserving im-
578 plementation of a DCOP algorithm that is based on lo-
579 cal search. It offers topology, constraint, and assign-
580 ment/decision privacy. The algorithm is much more scal-
581 able than P-Max-Sum, a privacy-preserving implementation
582 of another incomplete DCOP algorithm. It also offers so-
583 lutions with much better costs than those issued by P-Max-
584 Sum. Since P-DSA was able to solve in short time (3 min-
585 utes) problems involving as high as 100 agents, while all prior
586 studies on privacy-preserving DCOP algorithms report exper-
587 iments of much smaller scale and runtimes that are signifi-
588 cantly higher, P-DSA emerges as a suitable choice for solving
589 large-scale DCOPs in a privacy-preserving manner.

590 Our approach can also be extended to develop a privacy-
591 preserving version of MGM [Maheswaran *et al.*, 2004], an-
592 other local search algorithm for DCOPs. Even though the
593 “basic plot” in MGM is similar to DSA’s, it is more involved
594 as the decision to update local assignments is taken based on a
595 competition among agents and not by a coin-toss. Implement-
596 ing the more intricate logic of MGM in a privacy-preserving
597 manner is a challenge that we intend to undertake in a future
598 research.
599

548 **Number of agents in scale-free graphs.** We repeated the
549 previous experiment, but this time with scale-free graphs.
550 The results are given in Table 2. Here, too, we see that P-
551 DSA is more scalable and issues better solutions.

552 **Domain size in random graphs.** Here we fixed $n = 30$ and
553 $d = 0.4$ and varied the domain size m . The results are given
554 in Table 3. As already demonstrated, P-DSA is more scalable
555 than P-Max-Sum and managed to issue outputs to problems
556 in which P-Max-Sum failed to complete even one iteration
557 within the same time frame. Moreover, when both algorithms

600 Acknowledgments

601 This work was partially supported by the Ministry of Innova-
602 tion, Science and Technology, Israel.

603 References

- 604 [Barabási and Albert, 1999] A.L. Barabási and R. Albert.
605 Emergence of scaling in random networks. *Science*,
606 286:509–512, 1999.
- 607 [Burke and Brown, 2006] D. Burke and K. Brown. A com-
608 parison of approaches to handling complex local problems
609 in DCOP. In *Distributed Constraint Satisfaction Work-*
610 *shop*, pages 27–33, 2006.
- 611 [Chida *et al.*, 2018] K. Chida, D. Genkin, K. Hamada,
612 D. Ikarashi, R. Kikuchi, Y. Lindell, and A. Nof. Fast large-
613 scale honest-majority MPC for malicious adversaries. In
614 *CRYPTO*, pages 34–64, 2018.
- 615 [Damgård and Nielsen, 2007] I. Damgård and J.B. Nielsen.
616 Scalable and unconditionally secure multiparty computa-
617 tion. In *CRYPTO*, pages 572–590, 2007.
- 618 [Damle *et al.*, 2024] S. Damle, A. Triastcyn, B. Faltings, and
619 S. Gujar. Differentially private multi-agent constraint op-
620 timization. *Autonomous Agents and Multi-Agent Systems*,
621 38, 2024.
- 622 [Doshi *et al.*, 2008] P. Doshi, T. Matsui, M.C. Silaghi,
623 M. Yokoo, and M. Zanker. Distributed private constraint
624 optimization. In *WI-IAT*, pages 277–281, 2008.
- 625 [Faltings and Macho-Gonzalez, 2005] B. Faltings and
626 S. Macho-Gonzalez. Open constraint programming.
627 *Artificial Intelligence*, 161:1-2:181–208, 2005.
- 628 [Farinelli *et al.*, 2008] A. Farinelli, A. Rogers, A. Petcu, and
629 N.R. Jennings. Decentralised coordination of low-power
630 embedded devices using the max-sum algorithm. In *AA-*
631 *MAS*, pages 639–646, 2008.
- 632 [Fioretto *et al.*, 2016] F. Fioretto, W. Yeoh, and E. Pontelli.
633 Multi-variable agents decomposition for dcops. In *AAAI*,
634 volume 30, 2016.
- 635 [Fioretto *et al.*, 2017] F. Fioretto, W. Yeoh, and E. Pontelli.
636 A multiagent system approach to scheduling devices in
637 smart homes. In *AAAI workshops*, 2017.
- 638 [Gershman *et al.*, 2008] A. Gershman, A. Grubshtein,
639 A. Meisels, L. Rokach, and Roie Zivan. Scheduling
640 meetings by agents. In *PATAT*, 2008.
- 641 [Gershman *et al.*, 2009] A. Gershman, A. Meisels, and
642 R. Zivan. Asynchronous forward bounding for distributed
643 COPs. *Journal of Artificial Intelligence Research*, 34:61–
644 88, 2009.
- 645 [Greenstadt *et al.*, 2006] R. Greenstadt, J. Pearce, and
646 M. Tambe. Analysis of privacy loss in distributed con-
647 straint optimization. In *AAAI*, pages 647–653, 2006.
- 648 [Greenstadt *et al.*, 2007] R. Greenstadt, B. Grosz, and M.D.
649 Smith. SSDPOP: improving the privacy of DCOP with
650 secret sharing. In *AAMAS*, pages 171:1–171:3, 2007.

- [Grinshpoun and Tassa, 2016] T. Grinshpoun and T. Tassa. 651
P-SyncBB: A privacy preserving branch and bound DCOP 652
algorithm. *Journal of Artificial Intelligence Research*, 653
57:621–660, 2016. 654
- [Grinshpoun *et al.*, 2013] T. Grinshpoun, A. Grubshtein, 655
R. Zivan, A. Netzer, and A. Meisels. Asymmetric dis- 656
tributed constraint optimization problems. *Journal of Ar-* 657
tificial Intelligence Research, 47:613–647, 2013. 658
- [Grinshpoun *et al.*, 2019] T. Grinshpoun, T. Tassa, V. Levit, 659
and R. Zivan. Privacy preserving region optimal algo- 660
rithms for symmetric and asymmetric DCOPs. *Artificial* 661
Intelligence, 266:27–50, 2019. 662
- [Grinshpoun, 2012] T. Grinshpoun. When you say (DCOP) 663
privacy, what do you mean? - categorization of DCOP 664
privacy and insights on internal constraint privacy. In 665
ICAART, pages 380–386, 2012. 666
- [Grinshpoun, 2015] T. Grinshpoun. Clustering variables by 667
their agents. In *WI-IAT*, pages 250–256, 2015. 668
- [Hirayama and Yokoo, 1997] K. Hirayama and M. Yokoo. 669
Distributed partial constraint satisfaction problem. In *CP*, 670
pages 222–236, 1997. 671
- [Hirayama and Yokoo, 2005] K. Hirayama and M. Yokoo. 672
The distributed breakout algorithms. *Artificial Intelli-* 673
gence, 161:89–115, 2005. 674
- [Katagishi and Pearce, 2007] H. Katagishi and J.P. Pearce. 675
KOPT: Distributed DCOP algorithm for arbitrary k- 676
optima with monotonically increasing utility. In *DCR*, 677
2007. 678
- [Khakhiashvili *et al.*, 2021] I. Khakhiashvili, T. Grinshpoun, 679
and L. Dery. Course allocation with friendships as an 680
asymmetric distributed constraint optimization problem. 681
In *WI-IAT*, pages 688–693, 2021. 682
- [Kiekintveld *et al.*, 2010] C. Kiekintveld, Z. Yin, A. Kumar, 683
and M. Tambe. Asynchronous algorithms for approximate 684
distributed constraint optimization with quality bounds. In 685
AAMAS, pages 133–140, 2010. 686
- [Kim and Lesser, 2013] Y. Kim and V. Lesser. Improved 687
Max-Sum algorithm for DCOP with n-ary constraints. In 688
AAMAS, pages 191–198, 2013. 689
- [Kogan *et al.*, 2023] P. Kogan, T. Tassa, and T. Grinshpoun. 690
Privacy preserving solution of DCOPs by mediation. *Arti-* 691
ficial Intelligence, 319:103916, 2023. 692
- [Krigman *et al.*, 2024] S. Krigman, T. Grinshpoun, and 693
L. Dery. Scheduling of earth observing satellites using 694
distributed constraint optimization. *Journal of Scheduling*, 695
27:507–524, 2024. 696
- [Kumar *et al.*, 2008] A. Kumar, A. Petcu, and B. Faltings. H- 697
DPOP: Using hard constraints for search space pruning in 698
DCOP. In *AAAI*, pages 325–330, 2008. 699
- [Léauté and Faltings, 2013] T. Léauté and B. Faltings. Pro- 700
tecting privacy through distributed computation in multi- 701
agent decision making. *Journal of Artificial Intelligence* 702
Research, 47:649–695, 2013. 703

- 704 [Maheswaran *et al.*, 2004] R.T. Maheswaran, J.P. Pearce,
705 and M. Tambe. Distributed algorithms for DCOP: A
706 graphical-game-based approach. In *PDCS*, pages 432–
707 439, 2004.
- 708 [Maheswaran *et al.*, 2006] R.T. Maheswaran, J.P. Pearce,
709 E. Bowring, P. Varakantham, and M. Tambe. Privacy loss
710 in distributed constraint reasoning: A quantitative frame-
711 work for analysis and its applications. *Autonomous Agents
712 and Multi-Agent Systems*, 13:27–60, 2006.
- 713 [Modi *et al.*, 2005] P.J. Modi, W.M. Shen, M. Tambe, and
714 M. Yokoo. ADOPT: asynchronous distributed constraint
715 optimization with quality guarantees. *Artificial Intelli-
716 gence*, 161:149–180, 2005.
- 717 [Nguyen *et al.*, 2019] D.T. Nguyen, W. Yeoh, H.C. Lau, and
718 R. Zivan. Distributed gibbs: A linear-space sampling-
719 based DCOP algorithm. *Journal of Artificial Intelligence
720 Research*, 64:705–748, 2019.
- 721 [Nishide and Ohta, 2007] T. Nishide and K. Ohta. Multi-
722 party computation for interval, equality, and comparison
723 without bit-decomposition protocol. In *PKC*, pages 343–
724 360, 2007.
- 725 [Nissim and Zivan, 2005] K. Nissim and R. Zivan. Secure
726 DisCSP protocols - from centralized towards distributed
727 solutions. In *DCR Workshops*, 2005.
- 728 [Ottens *et al.*, 2017] B. Ottens, C. Dimitrakakis, and B. Fal-
729 tings. DUCT: An upper confidence bound approach to dis-
730 tributed constraint optimization problems. *ACM Transac-
731 tions on Intelligent Systems and Technology*, 8:69, 2017.
- 732 [Petcu and Faltings, 2005] A. Petcu and B. Faltings. A scal-
733 able method for multiagent constraint optimization. In *IJ-
734 CAI*, pages 266–271, 2005.
- 735 [Rust *et al.*, 2016] P. Rust, G. Picard, and F. Ramparany. Us-
736 ing message-passing DCOP algorithms to solve energy-
737 efficient smart environment configuration problems. In *IJ-
738 CAI*, pages 468–474, 2016.
- 739 [Savaux *et al.*, 2020] J. Savaux, J. Vion, S. Piechowiak,
740 R. Mandiau, T. Matsui, K. Hirayama, M. Yokoo, S. El-
741 mane, and M. Silaghi. Privacy stochastic games in dis-
742 tributed constraint reasoning. *Annals of Mathematics and
743 Artificial Intelligence*, 88:691–715, 2020.
- 744 [Shamir, 1979] A. Shamir. How to share a secret. *Communi-
745 cations of the ACM*, 22:612–613, 1979.
- 746 [Silaghi and Mitra, 2004] M.C. Silaghi and D. Mitra. Dis-
747 tributed constraint satisfaction and optimization with pri-
748 vacy enforcement. In *IAT*, pages 531–535, 2004.
- 749 [Silaghi *et al.*, 2006] M.C. Silaghi, B. Faltings, and A. Petcu.
750 Secure combinatorial optimization simulating DFS tree-
751 based variable elimination. In *AI&Math*, 2006.
- 752 [Tassa *et al.*, 2017] T. Tassa, T. Grinshpoun, and R. Zivan.
753 Privacy preserving implementation of the Max-Sum algo-
754 rithm and its variants. *Journal of Artificial Intelligence
755 Research*, 59:311–349, 2017.
- 756 [Tassa *et al.*, 2021] T. Tassa, T. Grinshpoun, and A. Yanai.
757 PC-SyncBB: A privacy preserving collusion secure DCOP
758 algorithm. *Artificial Intelligence*, 297:103501, 2021.
- [Teacy *et al.*, 2008] W.T.L. Teacy, A. Farinelli, N.J. Grab- 759
ham, P. Padhy, A. Rogers, and N.R. Jennings. Max-sum 760
decentralised coordination for sensor systems. In *AAMAS*, 761
pages 1697–1698, 2008. 762
- [Vion *et al.*, 2022] J. Vion, R. Mandiau, S. Piechowiak, and 763
M. Silaghi. Integrating domain and constraint privacy rea- 764
soning in the distributed stochastic algorithm with break- 765
outs. *Annals of Mathematics and Artificial Intelligence*, 766
90:31–73, 2022. 767
- [Yokoo and Hirayama, 2000] M. Yokoo and K. Hirayama. 768
Algorithms for distributed constraint satisfaction: A re- 769
view. *Autonomous Agents and Multi-Agent Systems*, 770
3:185–207, 2000. 771
- [Yokoo *et al.*, 2005] M. Yokoo, K. Suzuki, and K. Hirayama. 772
Secure distributed constraints satisfaction: Reaching 773
agreement without revealing private information. *Artifi- 774
cial Intelligence*, 161:229–246, 2005. 775
- [Zhang *et al.*, 2005] W. Zhang, G. Wang, Z. Xing, and 776
L. Wittenburg. Distributed stochastic search and dis- 777
tributed breakout: properties, comparison and applications 778
to constraint optimization problems in sensor networks. 779
Artificial Intelligence, 161:55–87, 2005. 780
- [Zivan *et al.*, 2014] R. Zivan, S. Okamoto, and H. Peled. Ex- 781
plorative anytime local search for distributed constraint 782
optimization. *Artificial Intelligence*, 212:1–26, 2014. 783