

GDBA Revisited: Unleashing the Power of Guided Local Search for Distributed Constraint Optimization

Yanchen Deng¹, Xinrun Wang^{2*}, Bo An¹

¹College of Computing and Data Science, Nanyang Technological University, Singapore

²School of Computing and Information Systems, Singapore Management University, Singapore

ycdeng@ntu.edu.sg, xrwang@smu.edu.sg, boan@ntu.edu.sg

Abstract

Local search is an important class of incomplete algorithms for solving Distributed Constraint Optimization Problems (DCOPs) but it often converges to poor local optima. While Generalized Distributed Breakout Algorithm (GDBA) provides a comprehensive rule set to escape premature convergence, its empirical benefits remain marginal on general-valued problems. In this work, we systematically examine GDBA and identify three factors that potentially lead to its inferior performance, i.e., over-aggressive constraint violation conditions, unbounded penalty accumulation, and uncoordinated penalty updates. To address these issues, we propose Distributed Guided Local Search (DGLS), a novel GLS framework for DCOPs that incorporates an adaptive violation condition to selectively penalize constraints with high cost, a penalty evaporation mechanism to control the magnitude of penalization, and a synchronization scheme for coordinated penalty updates. We theoretically show that the penalty values are bounded, and agents play a potential game in DGLS. Extensive empirical results on various benchmarks demonstrate the great superiority of DGLS over state-of-the-art baselines. Compared to Damped Max-sum with high damping factors, our DGLS achieves competitive performance on general-valued problems, and outperforms by significant margins on structured problems in terms of anytime results.

Code — <https://github.com/ycdeng-ntu/DGLS>

Extended version — <https://arxiv.org/pdf/2508.06899>

1 Introduction

Distributed Constraint Optimization Problems (DCOPs) (Modi et al. 2005; Fioretto, Pontelli, and Yeoh 2018) are a fundamental formalism for cooperative multi-agent systems where a set of autonomous agents coordinate with each other to pursue a global objective via localized communication. DCOPs have been successfully applied to model various real-world applications, including scheduling (Pertzovsky, Zivan, and Agmon 2024), resource allocation (Monteiro et al. 2012), and smart grids (Fioretto et al. 2017).

Efficiently solving DCOPs remains a long-standing challenge due to the inherent NP-hardness. Over the past

decades, many DCOP algorithms have been proposed and are generally categorized into *complete* algorithms and *incomplete* algorithms, according to whether they guarantee finding the optimal solutions. Complete algorithms include distributed backtracking search (Modi et al. 2005; Checheta and Sycara 2006; Yeoh, Felner, and Koenig 2010; Netzer, Grubshtein, and Meisels 2012; Litov and Meisels 2017; Deng et al. 2019) and inference (Petcu and Faltings 2005, 2007; Chen et al. 2020), which exhaust the solution space by either branch-and-bound or bucket elimination (Dechter 1998). However, the coordination overheads of these algorithms scale exponentially w.r.t. the problem size, making them unsuitable for large-scale applications.

On the other hand, incomplete algorithms trade the optimality for practical computational overheads (Farinelli et al. 2008; Cohen, Galiki, and Zivan 2020; Chen et al. 2018; Nguyen et al. 2019; Ottens, Dimitrakakis, and Faltings 2017). Among them, local search (Zhang et al. 2005; Maheswaran, Pearce, and Tambe 2004; Pearce and Tambe 2007; Zivan, Okamoto, and Peled 2014; Hoang et al. 2018) is an important class of incomplete algorithms, which iteratively refines the solution via local moves. However, these algorithms often prematurely converge to poor local optima due to their greedy nature. As an instantiation of Guided Local Search (GLS) (Voudouris, Tsang, and Alsheddy 2010), Generalized Distributed Breakout Algorithm (GDBA) (Okamoto, Zivan, and Nahon 2016) was introduced to provide a comprehensive rule set for breaking out of the local optima by adapting the notion of constraint violation and cost increase in DBA (Hirayama and Yokoo 2005). However, the empirical benefits of GDBA often appear to be marginal compared to well-established baselines like DSA (Zhang et al. 2005) on general-valued problems.

To understand why GDBA underperforms on general-valued problems, we first conduct a pilot study (cf. Figure 1) analyzing the penalty dynamics of GDBA on both random DCOPs and structured problems. Our observations reveal that, on general-valued instances, GDBA uniformly accumulates penalties for nearly all constraints, making them receive heavy but similar attention, which offsets the benefit of breakout. Such ineffective penalization stems from over-aggressive constraint violation conditions (e.g., NM) that classify most constraints as violated, monotonic penalty increases that lead to unbounded penalty accumulation, and

*Corresponding author.

uncoordinated penalty updates that cause agents to optimize misaligned objective functions. In light of this, we present a novel Distributed Guided Local Search (DGLS) framework for DCOPs, which incorporates an adaptive constraint violation condition based on the costs of each constraint, an evaporation mechanism to avoid unbounded penalty accumulation, and a synchronization scheme to enforce coordinated penalty updates. Specifically, our contributions are:

- We systematically examine the key design choices of GDBA. We find that the over-aggressive constraint violation conditions, monotonic penalty increment and uncoordinated penalty update would lead to inferior performance on general-valued DCOPs.
- We present a novel DGLS framework that enables efficient GLS for DCOPs. Our DGLS incorporates an adaptive constraint violation condition, an evaporation mechanism, and a penalty synchronization scheme to fully unleash the performance of GLS in solving DCOPs. We also theoretically show that the penalty values are bounded, and agents play a potential game in our DGLS.
- We compare our DGLS with state-of-the-art DCOP algorithms on various standard benchmarks. Our extensive empirical results show the great potential of DGLS: on general-valued problems and scale-free network, DGLS is able to match or perform slightly better than Damped Max-sum (DMS) (Cohen, Galiki, and Zivan 2020) with a high damping factor, while DGLS outperforms DMS by significant margins on structured problems including 2D lattices, meeting scheduling and weighted graph coloring in terms of anytime performance (Zilberstein 1996; Zivan, Okamoto, and Peled 2014).

2 Preliminaries

In this section, we review necessary preliminaries including DCOPs, GLS and GDBA.

2.1 Distributed Constraint Optimization Problems

A DCOP (Modi et al. 2005) can be formalized as a tuple $\langle I, X, D, F \rangle$ where $I = \{1, \dots, |I|\}$ is the set of agents, $X = \{x_1, \dots, x_{|X|}\}$ is the set of variables, $D = \{D_1, \dots, D_{|X|}\}$ is the set of discrete domains, and $F = \{f_1, \dots, f_{|F|}\}$ is the set of constraint functions. Each variable x_i takes a value from its domain D_i , and each constraint function $f_i : D_{i_1} \times \dots \times D_{i_k} \rightarrow \mathbb{R}_{\geq 0}$ defines a cost for each possible combination of variables $scp(f_i) = (x_{i_1}, \dots, x_{i_k})$. The objective is to find a solution $\tau^* = (d_1^*, \dots, d_{|X|}^*)$ such that the total cost is minimized:

$$\tau^* = \arg \min_{\tau \in \prod_i D_i} \sum_{f_j \in F} f_j(\tau|_{scp(f_j)}), \quad (1)$$

where $\tau|_{scp(f_j)}$ is the projection of τ onto $scp(f_j) \subseteq X$. For the sake of simplicity, we follow the common assumptions that each agent controls only one variable (i.e., $|I| = |X|$) and all constraints are binary (i.e., $f_{ij} : D_i \times D_j \rightarrow \mathbb{R}_{\geq 0}, \forall f_{ij} \in F$). Therefore, the terms ‘‘agent’’ and ‘‘variable’’ can be used interchangeably.

2.2 Guided Local Search

GLS (Voudouris, Tsang, and Alsheddy 2010) is a meta-heuristic for helping local search algorithms to escape local optima using a penalty system. Specifically, GLS considers the following augmented objective function:

$$h(\tau) = f(\tau) + \lambda \sum_i p_i \cdot \mathbb{I}[\text{feature } i \text{ presents in } \tau], \quad (2)$$

where $f(\cdot)$ is the original objective function, p_i is the penalty associated with feature i , \mathbb{I} is the indicator function, and $\lambda > 0$ is the weight to balance the original objective and penalty. Here, the features are problem-specific. For example, a feature could be an edge from city A to city B in a Traveling Salesman Problem (TSP), or whether a hard clause is satisfied in Boolean Satisfiability (SAT) (Cai and Lei 2020). When local search converges to a local optimum, GLS will select a subset of features presented in the incumbent solution to increase the associated penalty value, so as to force local search to explore novel solutions. In other words, GLS breaks out of the local optimum by modifying the problem’s objective landscape.

2.3 Generalized Distributed Breakout Algorithm

GDBA (Okamoto, Zivan, and Nahon 2016) provides a comprehensive set of rules for breaking out of local optima in DCOPs, which adapts the concepts of constraint violation and cost increase in DBA (Hirayama and Yokoo 2005) for solving Distributed Constraint Satisfaction Problems (DisCSPs) (Yokoo et al. 1998). Essentially, GDBA can be viewed as an instantiation of GLS where the features can be a full constraint function, a specific row or column within it, or a single cost cell, depending on the scope of changes to the penalty values during breakouts. Specifically, when a Quasi Local Minimum (QLM) (Hirayama and Yokoo 2005) is detected, each agent in QLM first identifies a set of violated constraints based on the cost values under the incumbent local solution, then *independently* increases the penalty associated with the features in these violated constraints by 1, which corresponds to selective penalization in GLS. Besides additive penalty like Eq. (2), GDBA’s variants also consider the multiplicative penalty:

$$h(\tau) = \sum_{f_{ij} \in F} f_{ij}(d_i, d_j) \cdot [1 + M_{ij}(d_i, d_j)], \quad (3)$$

where $\tau = (d_1, \dots, d_{|X|})$ is a solution, M_{ij} is the cost modifier (i.e., a matrix of penalty values) associated with constraint function f_{ij} with the initial value of 0.

3 Distributed Guided Local Search

In this section, we present the Distributed Guided Local Search (DGLS) framework. We first motivate our research by analyzing the key design choices of GDBA. Then we detail DGLS and theoretically analyze its properties.

3.1 Motivation

To analyze the dynamics of penalty values of GDBA, we consider its widely used variant $\langle M, NM, T \rangle$ on both sparse

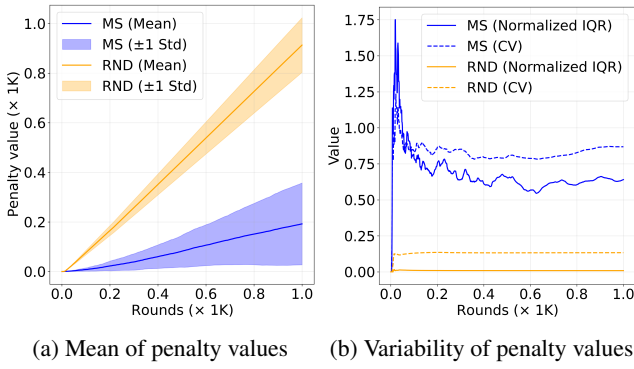


Figure 1: Analysis of penalty dynamics of GDBA

random DCOPs (RND) and meeting scheduling (MS) (Zivan, Okamoto, and Peled 2014; Maheswaran et al. 2004) where GDBA demonstrates notably inferior and superior performance, respectively (cf. Section 4). Figure 1a plots the mean penalty values in all cost modifiers against iterations.

It can be seen that the average penalty grows linearly in random DCOPs, meaning that almost every constraint is penalized in most rounds. That is due to the fact that the non-minimum (NM) violation condition is over-aggressive for general-valued constraints where the minimum cost entries are sparse. As a result, most constraints are considered as violated and the corresponding cost modifiers are increased once a QLM is detected. In contrast, when solving meeting scheduling problem where multiple minimum costs present in each constraint, penalty growth is much more moderate.

To look deeper into this issue, we also display the normalized IQR (i.e., interquartile range divided by mean) and Coefficient of Variation (CV) of penalty values in Figure 1b. It can be observed that, compared to solving meeting scheduling problems, penalty values in GDBA have a much lower variability when solving random DCOPs. Such indiscriminate penalization cause each constraint to receive heavy but similar attention, which could offset the benefit of breakout.

The unbounded penalty accumulation also contributes to the ineffective penalization. Specifically, GDBA monotonically increases cost modifiers without any decay mechanism, even when the constraints are almost perfectly satisfied, e.g., constraint costs that are only slightly above the minimum costs, which in turn exacerbates the pathologies of over-penalization and indiscriminate penalization, as evidenced by Figure 1a and 1b, respectively.

Finally, agents in QLM independently increase the cost modifiers, which may cause agents to optimize different objectives. Consider a constraint f_{ij} being flagged as violated for the first time. If agent i is in QLM while agent j doesn't, then the cost modifier for f_{ij} becomes 1 from agent i 's side but remains 0 from agent j 's side after breakout. Such mismatch introduces asymmetry in cost structures and potentially breaks the correspondence between pure strategy Nash Equilibrium and local optimum (Chapman et al. 2011; Grinshpoun et al. 2013).

Algorithm 1: Distributed Guided Local Search for agent i

```

1: Initialize cost modifiers to 0
2: Choose a random assignment  $d_i \in D_i$ 
3: Send  $x_i = d_i$  to all neighbors  $\mathcal{N}_i$ 
4: while termination condition is not met do
5:    $\bar{P}_i \leftarrow \emptyset$ 
6:   Receive assignment  $x_j = d_j$  from all neighbors  $\mathcal{N}_i$ 
7:    $d_i^* \leftarrow \arg \min_{d \in D_i} \sum_j \text{EFFCOST}(d, j, d_j)$ 
8:    $\Delta_i \leftarrow \sum_j \text{EFFCOST}(d_i, j, d_j) - \text{EFFCOST}(d_i^*, j, d_j)$ 
9:   Send gain  $\Delta_i$  to all neighbors  $\mathcal{N}_i$ 
10:  Receive gain  $\Delta_j$  from all neighbors  $\mathcal{N}_i$ 
11:  if  $\Delta_i > 0$  then
12:    if  $\Delta_i$  is the best improvement then
13:       $d_i \leftarrow d_i^*$ 
14:    else if no neighbor can improve then
15:      for  $j \in \mathcal{N}_i$  do
16:        if  $\text{ISVIOLATED}(d_i, j, d_j)$  then
17:           $\bar{P}_i \leftarrow \bar{P}_i \cup \{j\}$ 
18:          Send a (SYNC,  $i$ ) message to agent  $j$ 
19:         $\tilde{P}_i \leftarrow$  receive (SYNC,  $j$ ) from neighbors  $\mathcal{N}_i$ 
20:      for  $j \in \mathcal{N}_i$  do
21:         $\text{EVAPORATE}(j)$ 
22:         $\text{INCREASEMOD}(d_i, j, d_j, \bar{P}_i, \tilde{P}_i)$ 
23:      Send  $x_i = d_i$  to all neighbors  $\mathcal{N}_i$ 

```

Algorithm 2: Effective Cost Computation

```

1: function  $\text{EFFCOST}(d_i, j, d_j)$ 
2:   if  $\text{manner} = \text{additive}$  then
3:     return  $f_{ij}(d_i, d_j) + M_{ij}(d_i, d_j)$ 
4:   else if  $\text{manner} = \text{multiplicative}$  then
5:     return  $f_{ij}(d_i, d_j) \cdot [M_{ij}(d_i, d_j) + 1]$ 

```

3.2 DGLS Framework

In light of this, we present a novel Distributed Guided Local Search (DGLS) framework for DCOPs. To overcome the over-penalization and indiscriminate penalization, our DGLS incorporates an adaptive rule for selectively penalizing the constraints with high costs. Besides, we introduce an evaporation mechanism to avoid unbounded penalty accumulation. Finally, we propose a synchronization scheme to enforce coordinated penalty update. Algorithm 1 presents the sketch of DGLS.

Like GDBA, each agent i in our DGLS begins with initializing the cost modifiers and broadcasting a randomly selected assignment d_i to its neighbors. After that, it looks for the best assignment d_i^* and calculates the corresponding gain Δ_i based on the current local view and cost modifiers via EFFCOST (cf. Algorithm 2), where the base cost is modified in either additive (A) or multiplicative (M) way. After that, it broadcasts the gain to the neighbors to determine whether it should make a local move. Particularly, if $\Delta_i = 0$ and no neighbor can improve, then a QLM is detected and agent i starts to select constraints to penalize. After that, all cost modifiers are evaporated by a rate γ , and finally agent i performs coordinated penalty update given a scope (i.e., cell, table, row or column). Therefore, a DGLS is instantiated by specifying a tuple $(A/M, \gamma, \text{cel}/\text{tab}/\text{row}/\text{col})$.

Algorithm 3: Adaptive Violation Condition

```

1: function ISVIOLATED( $d_i, j, d_j$ )
2:    $\eta \leftarrow \frac{f_{ij}(d_i, d_j) - \check{f}_{ij}}{\check{f}_{ij} - \hat{f}_{ij}}$ 
3:   if RANDOM( $0, 1$ )  $< \eta$  then
4:     return True
5:   return False

```

Adaptive violation condition Instead of using fixed rules in GDBA, our DGLS incorporates an adaptive rule for identifying violated constraints (cf. Algorithm 3). Intuitively, for each constraint f_{ij} we first compute its normalized cost η , where $\check{f}_{ij} = \min_{d_i, d_j} f_{ij}(d_i, d_j)$ and $\hat{f}_{ij} = \max_{d_i, d_j} f_{ij}(d_i, d_j)$. Then we stochastically mark the constraint as violated with a probability of η . This also generalizes the utility-based penalization in classical GLS (Voudouris, Tsang, and Alsheddy 2010), where the features with the highest score are deterministically penalized.

Our adaptive rule has several unique advantages. First, it eliminates the need to tune the constraint violation condition in GDBA (e.g., *NZ*, *NM* and *MX*). Second, it perfectly aligns with objective (1) by measuring the “badness” of a constraint with the normalized cost. Particularly, if the constraint cost equals the minimum value, then the constraint cannot be flagged as violated; conversely, if the constraint cost attains the maximum value, then there must be a constraint violation. By selectively penalizing constraints, we avoid over-penalizing the constraints with the cost close to their minimum value, which directs local search to pay more attention to the constraints that incur high costs.

Evaporation mechanism We periodically decay the cost modifiers through an evaporation mechanism (cf. EVAPORATE). For each constraint f_{ij} , the associated cost modifier M_{ij} is geometrically decayed by $0 < \gamma < 1$. Formally,

$$M_{ij}(d_i, d_j) \leftarrow \gamma M_{ij}(d_i, d_j), \quad \forall d_i \in D_i, d_j \in D_j. \quad (4)$$

The evaporation mechanism addresses the issue of unbounded penalty accumulation. Combining with our adaptive violation condition, it helps local search to effectively forget the penalty of the well-satisfied constraints (e.g., the constraints whose current cost is close to the minimum value), and thus realizes selective penalization.

Coordinated penalty update We coordinate the penalty update between two agents by explicitly communicating the index of constraints to be penalized. Specifically, for each round, agent i maintains a set of self-penalized constraints \bar{P}_i and a set of constraints penalized by neighbors \tilde{P}_i . When it decides to penalize a constraint f_{ij} (i.e., ISVIOLATED(d_i, j, d_j)=True), agent i first records the index to \bar{P}_i and notify neighbor j through a SYNC message. After that, it collects the index associated with all SYNC messages from neighbors as \tilde{P}_i . Finally, it performs coordinated penalty update according to Algorithm 4.

Specifically, if the update scope is either cell (*cel*) or table (*tab*), then the corresponding entries of the cost modifier will be increased by 1 if the index presents in either \bar{P}_i or

Algorithm 4: Coordinated Penalty Update

```

1: function INCREASEMOD( $d_i, j, d_j, \bar{P}_i, \tilde{P}_i$ )
2:   if scope = cell then
3:     if  $j \in \bar{P}_i \vee j \in \tilde{P}_i$  then
4:        $M_{ij}(d_i, d_j) \leftarrow M_{ij}(d_i, d_j) + 1$ 
5:   else if scope = table then
6:     if  $j \in \bar{P}_i \vee j \in \tilde{P}_i$  then
7:        $M_{ij}(d'_i, d'_j) \leftarrow M_{ij}(d'_i, d'_j) + 1, \forall d'_i, d'_j$ 
8:   else if scope = row then
9:     if  $j \in \bar{P}_i$  then
10:       $M_{ij}(d_i, d'_j) \leftarrow M_{ij}(d_i, d'_j) + 1, \forall d'_j$ 
11:    if  $j \in \tilde{P}_i$  then
12:       $M_{ij}(d'_i, d_j) \leftarrow M_{ij}(d'_i, d_j) + 1, \forall d'_i$ 
13:    if  $j \in \bar{P}_i \wedge j \in \tilde{P}_i$  then
14:       $M_{ij}(d_i, d_j) \leftarrow M_{ij}(d_i, d_j) - 1$ 
15:   else if scope = column then
16:     if  $j \in \bar{P}_i$  then
17:        $M_{ij}(d'_i, d_j) \leftarrow M_{ij}(d'_i, d_j) + 1, \forall d'_i$ 
18:     if  $j \in \tilde{P}_i$  then
19:        $M_{ij}(d_i, d'_j) \leftarrow M_{ij}(d_i, d'_j) + 1, \forall d'_j$ 
20:     if  $j \in \bar{P}_i \wedge j \in \tilde{P}_i$  then
21:        $M_{ij}(d_i, d_j) \leftarrow M_{ij}(d_i, d_j) - 1$ 

```

\tilde{P}_i . On the other hand, if the update scope is row, agent i will increase the entries in the d_i -th row regardless of x_j 's assignment if the penalization of f_{ij} is initiated by agent i . If $j \in \tilde{P}_i$, agent i mirrors the operation of agent j by penalizing the d_j -th column of the cost modifier. Finally, if both agent i and j penalize constraint f_{ij} , then we minus 1 from $M_{ij}(d_i, d_j)$ to avoid double-counting. The scope of column (*col*) follows a similar pattern by exchanging row and column operations.

3.3 Theoretical Results

In this subsection, we theoretically analyze the properties of DGLS. We first show the upper bound of values in cost modifiers in our DGLS, which avoids the uncontrollably growing penalty values in GDBA.

Theorem 1. *With evaporation rate $0 < \gamma < 1$, the penalty values in any cost modifier are bounded by $1/(1 - \gamma)$.*

Proof. Consider the worst-case scenario where a specific entry (d_i, d_j) is incremented in every round. Let $M_{ij}^{(k)}$ be the penalty value in round k . We have

$$\begin{aligned}
M_{ij}^{(1)}(d_i, d_j) &= 0 \cdot \gamma + 1 = 1 \\
M_{ij}^{(2)}(d_i, d_j) &= 1 \cdot \gamma + 1 = 1 + \gamma \\
&\dots\dots \\
M_{ij}^{(k)}(d_i, d_j) &= 1 + \gamma + \gamma^2 + \dots + \gamma^{k-1}.
\end{aligned}$$

As $k \rightarrow \infty$, this geometric series converges to $1/(1 - \gamma)$, which concludes the theorem. \square

Corollary 1. *The effective cost is bounded for both additive and multiplicative cases:*

$$\begin{aligned}\text{EFFCOST}_A(d_i, j, d_j) &\leq \hat{f}_{ij} + 1/(1 - \gamma) \\ \text{EFFCOST}_M(d_i, j, d_j) &\leq \hat{f}_{ij} \cdot [1 + 1/(1 - \gamma)].\end{aligned}$$

This corollary indicates that the contribution of any constraint to the total effective cost is bounded by a constant, which prevents any constraint from dominating the augmented objective as the number of rounds grows.

We then show that our coordinated penalty update guarantees the consistency of cost modifiers from both sides of each constraint. This consistency, in turn, enables a potential game structure in DGLS.

Lemma 1. *At the beginning of each round, the cost modifier of the constraint between agent i and agent j from i 's side is the same as the counterpart from j 's side.¹*

Theorem 2. *For each round, agents in DGLS play a potential game where the potential function is the total effective cost given the current cost modifiers.*

Proof. Define the potential function as

$$\Phi(\tau) = \frac{1}{2} \sum_{i \in I} \sum_{j \in \mathcal{N}_i} \text{EFFCOST}(d_i, j, d_j), \quad (5)$$

where $\tau = (d_1, \dots, d_{|X|})$ is a solution. Consider an agent i unilaterally changing its assignment from d_i to d'_i , while all other agents maintain their current assignments. The change in agent i 's local cost is:

$$\Delta_i = \sum_{j \in \mathcal{N}_i} \text{EFFCOST}(d_i, j, d_j) - \text{EFFCOST}(d'_i, j, d_j).$$

The change in the potential function is:

$$\begin{aligned}\Delta\Phi &= \frac{1}{2} \sum_{j \in \mathcal{N}_i} \text{EFFCOST}(d_i, j, d_j) - \text{EFFCOST}(d'_i, j, d_j) \\ &\quad + \frac{1}{2} \sum_{j \in \mathcal{N}_i} \text{EFFCOST}(d_j, i, d_i) - \text{EFFCOST}(d_j, i, d'_i).\end{aligned}$$

By Lemma 1, the cost modifiers are the same from both agent i 's and j 's side, which implies

$$\text{EFFCOST}(d_i, j, d_j) = \text{EFFCOST}(d_j, i, d_i),$$

and

$$\text{EFFCOST}(d'_i, j, d_j) = \text{EFFCOST}(d_j, i, d'_i).$$

Therefore, $\Delta_i = \Delta\Phi$, which concludes the theorem. \square

Theorem 2 indicates that unlike in GDBA, agents in our DGLS optimize a consistent and well-defined global objective, i.e., the total effective cost (cf. Eq. (2) and Eq. (3)). Furthermore, any local improvement of an agent exactly corresponds to a reduction in the global effective cost.

We now establish some equivalences of DGLS variants.

Theorem 3. *If $f_{ij} : D_i \times D_j \rightarrow \{0, 1\}$, $\forall f_{ij} \in F$, then (A, γ, cel) and (M, γ, cel) are equivalent for any γ .*

¹Proof is provided in the extended version.

Proof. We prove the theorem by showing that for each round k and assignments $d_i \in D_i, d_j \in D_j$,

$$\text{EFFCOST}_A^{(k)}(d_i, j, d_j) = \text{EFFCOST}_M^{(k)}(d_i, j, d_j), \quad (6)$$

where $\text{EFFCOST}_A^{(k)}$ and $\text{EFFCOST}_M^{(k)}$ are the effective cost for round k under additive and multiplicative manner, respectively. If $f_{ij}(d_i, d_j) = 0$, then the constraint f_{ij} cannot be flagged as violated since $\eta = 0$ (cf. Algorithm 3). Therefore the corresponding entry (d_i, d_j) of the cost modifier M_{ij} in both variants remains 0 since no increment happens. In this case, Eq. (6) holds because

$$\text{EFFCOST}_A^{(k)}(d_i, j, d_j) = 0 + M_{ij}(d_i, d_j) = 0 + 0 = 0$$

$$\text{EFFCOST}_M^{(k)}(d_i, j, d_j) = 0 \cdot [1 + M_{ij}(d_i, d_j)] = 0 \cdot 1 = 0.$$

On the other hand, if $f_{ij}(d_i, d_j) = 1$, then the two variants maintain the same penalty for entry (d_i, d_j) given the same decay factor γ . We show it by induction. The fact holds for the first round where the cost modifiers are 0. Assume that it holds for round k . If there is a QLM and d_i, d_j are the incumbent assignments, then f_{ij} must be violated and the penalty of entry (d_i, d_j) will be incremented by 1 in both variants since $\eta = 1$. Otherwise, there is no increment for both variants. Given the same evaporation rate γ , the fact holds for round $k + 1$. Therefore, Eq. (6) holds for this case because

$$\text{EFFCOST}_A^{(k)}(d_i, j, d_j) = 1 + M_{ij}(d_i, d_j)$$

$$\text{EFFCOST}_M^{(k)}(d_i, j, d_j) = 1 \cdot [1 + M_{ij}(d_i, d_j)] = 1 + M_{ij}(d_i, d_j). \quad \square$$

Theorem 4. *(A, γ, tab) is equivalent to Maximum Gain Message (MGM) for any γ .*

Proof. Let's consider the decision process of agent i under the additive manner:

$$\begin{aligned}d_i^* &= \arg \min_{d'_i \in D_i} \sum_{j \in \mathcal{N}_i} \text{EFFCOST}(d'_i, j, d_j) \\ &= \arg \min_{d'_i \in D_i} \sum_{j \in \mathcal{N}_i} f_{ij}(d'_i, d_j) + M_{ij}(d'_i, d_j).\end{aligned}$$

In table scope, the cost modifiers are updated table-wise (cf. line 5-7 of Algorithm 4). Therefore $M_{ij}(d'_i, d_j) = c_{ij}, \forall d'_i \in D_i$, which gives us

$$\begin{aligned}d_i^* &= \arg \min_{d'_i \in D_i} \sum_{j \in \mathcal{N}_i} f_{ij}(d'_i, d_j) + c_{ij} \\ &= \arg \min_{d'_i \in D_i} \left(c + \sum_{j \in \mathcal{N}_i} f_{ij}(d'_i, d_j) \right) \\ &= \arg \min_{d'_i \in D_i} \sum_{j \in \mathcal{N}_i} f_{ij}(d'_i, d_j).\end{aligned}$$

This exactly matches the decision rule for agents in MGM (Maheswaran, Pearce, and Tambe 2004). Besides,

$$\begin{aligned}\Delta_i &= \sum_{j \in \mathcal{N}_i} \text{EFFCOST}(d_i, j, d_j) - \text{EFFCOST}(d_i^*, j, d_j) \\ &= \sum_{j \in \mathcal{N}_i} f_{ij}(d_i, d_j) + M_{ij}(d_i, d_j) - f_{ij}(d_i^*, d_j) - M_{ij}(d_i^*, d_j) \\ &= \sum_{j \in \mathcal{N}_i} f_{ij}(d_i, d_j) - f_{ij}(d_i^*, d_j) + c_{ij} - c_{ij} \\ &= \sum_{j \in \mathcal{N}_i} f_{ij}(d_i, d_j) - f_{ij}(d_i^*, d_j),\end{aligned}$$

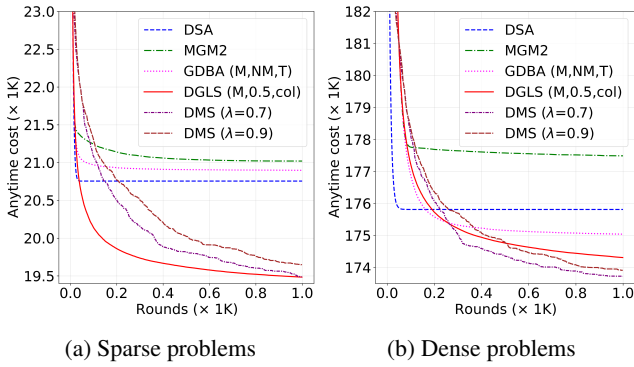


Figure 2: Performance on random DCOPs

which also produces the same gain as in MGM. Therefore, the theorem is concluded. \square

Finally, we show the complexity of DGLS as follows.

Theorem 5. *In each round of DGLS, agent i communicates $O(|\mathcal{N}_i|)$ messages and performs $O(|\mathcal{N}_i| * |D_{\max}^i| * |D_i|)$ operations, where $D_{\max}^i = \arg \max_{j \in \mathcal{N}_i} |D_j|$.*

Proof. Each round agent i communicates $|\mathcal{N}_i|$ assignment messages, $|\mathcal{N}_i|$ gain messages, and q SYNC messages, where $q \leq |\mathcal{N}_i|$ since i notifies a neighbor j only when f_{ij} is flagged as violated. Therefore, the total number of messages communicated by agent i is in $O(|\mathcal{N}_i|)$.

Besides, agent i finds d_i^* in $O(|\mathcal{N}_i| * |D_i|)$ operations, evaporates all cost modifiers in $O(|\mathcal{N}_i| * |D_{\max}^i| * |D_i|)$ operations and increments the cost modifiers in $O(|\mathcal{N}_i| * |D_{\max}^i| * |D_i|)$ operations in the worst case (e.g., *tab* scope is used and all involved constraints are penalized). Therefore, the total number of operations is in $O(|\mathcal{N}_i| * |D_{\max}^i| * |D_i|)$. \square

4 Empirical Evaluations

Benchmarks and baselines We evaluate our algorithm on various standard DCOP benchmarks, including: (1) random DCOPs with 120 agents, a graph density of 0.1 (sparse) or 0.6 (dense); (2) scale-free networks (Barabási and Albert 1999) with 120 agents and $m_0 = m_1 = 3$; (3) 2D lattices with grid size of 10×10 ; (4) meeting scheduling problems using Events-as-Variable (EAV) formulation (Zivan, Okamoto, and Peled 2014; Maheswaran et al. 2004) with 20 available time slots, 20 meetings and 90 persons, where each person randomly selects two meetings to participate, the travel time between any pair of meetings is uniformly drawn from $[6, 10]$, and a cost equal to the number of overlooked persons is incurred if the difference between the time slots of two meetings are less than the travel time; (5) weighted graph coloring problems with 120 agents, 3 available colors, and graph density of 0.05, where a conflict cost is uniformly selected from $[1, 100]$ if two adjacent agents assign the same color. For benchmarks (1–3), we consider the problems with a domain size of 10 and uniformly select constraint costs from $[0, 100]$. For each benchmark, we generate




Figure 3: Performance on topology-structured problems

100 random instances and each instance is solved 20 times with the maximum round of 1000. Finally, we average the anytime cost (Zilberstein 1996; Zivan, Okamoto, and Peled 2014) of each round of 2000 runs as the final result. All experiments are conducted on a Linux workstation with Intel Xeon W-2133 CPU and 32GB memory.

For competitors, we consider DSA (Zhang et al. 2005) with $p = 0.8$ and GDBA (Okamoto, Zivan, and Nahon 2016) as representative local search algorithms; MGM2 (Maheswaran, Pearce, and Tambe 2004) as a representative K -OPT algorithm, and DMS (Cohen, Galiki, and Zivan 2020) with damping factors $\lambda = 0.7$ and 0.9 as the state-of-the-art baselines. For GDBA, we consider its (M, NM, T) variant that exhibits the strongest performance according to (Okamoto, Zivan, and Nahon 2016).

Performance comparison Empirically, we find DGLS variants $(M, 0.5, col)$ and $(M, 0.9, col)$ perform best on general-valued problems and cost-structured problems, respectively. Therefore, we consider these two variants for performance comparison. Figure 2 presents the anytime results on both sparse and dense random DCOPs. It can be observed that our DGLS exhibits a fast convergence speed on sparse problems, quickly surpasses all baselines after about 50 rounds. GDBA, on the other hand, fails to effectively break out of local optima and is dominated by DSA in the sparse case. Besides, it is interesting to find that the gap between GDBA and our DGLS narrows on dense problems. This is because each constraint has a higher cost than the sparse setting. Therefore, agents in our DGLS trigger penalization more frequently (cf. Algorithm 3), which exhibits similar behavior to the GDBA under the NM violation condition. Still, DGLS outperforms GDBA by a significant margin in the dense setting, thanks to the evaporation mechanism and synchronization scheme for coordinated penalty updates.

Figure 3 compares the anytime performance on topology-structured problems. It can be seen that GDBA performs poorly and is strictly dominated by all other competitors on both scale-free networks and 2D lattices, which highlights the inefficiency of GDBA in dealing with general-valued problems. In contrast, our DGLS effectively changes the problems' landscape through adaptive violation condition,

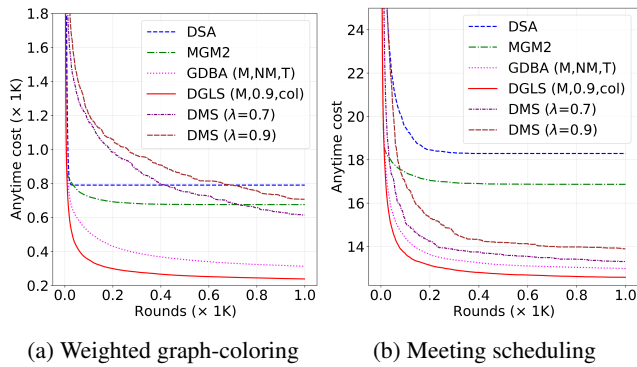


Figure 4: Performance on cost-structured problems

evaporation and coordinated penalty update whenever local search gets trapped in local optima. These mechanisms enable selective penalization that directs local search to focus more on the constraints with high costs, and therefore result in a much steadier improvement over time. In fact, DGLS not only significantly outperforms all local search heuristics, but also improves DMS by 3.77%–6.03% on 2D lattices with p -value $< 10^{-5}$.

Figure 4 presents the results on weighted graph-coloring problems and meeting scheduling problems. GDBA demonstrates great advantages over all competitors except DGLS on these cost-structured problems. This can be attributed to the multiple cost minimums present in each constraint (e.g., the cost entries for conflict-free colors in weighted graph coloring problems or time slots in meeting scheduling problems). As a consequence, the cost modifiers in GDBA grow moderately compared to general-valued problems (cf. Figure 1), thus the over-penalization and indiscriminate penalization are alleviated as a side effect. Nevertheless, GDBA is still strictly dominated by DGLS which explicitly implements selective penalization through the adaptive violation condition and evaporation mechanism. Notably, our DGLS surpasses all baselines within the first 50 rounds and continuously improves given more rounds, significantly outperforming DMS by 61.24%–66.30% and 5.47%–9.45% on weighted graph-coloring problems and meeting scheduling problems with p -value $< 10^{-5}$, respectively.

Ablation study To understand how each component contributes to the success of our DGLS, we perform an ablation study on sparse random DCOPs and present results in Figure 5. Here, we consider the DGLS variant of $(M, 0.5, tab)$ for ablation since it is algorithmically comparable to GDBA (M, NM, T) due to the same manner and scope parameters. In fact, GDBA (M, NM, T) can be recovered from DGLS $(M, 0.5, tab)$ if adaptive violation condition (AVC), evaporation, and coordinated penalty update (CPU) are disabled.

AVC tends to contribute most significantly to the anytime performance. Without AVC, DGLS indiscriminately penalizes each constraint that incurs a cost higher than the minimum cost, which leads to ineffective penalization and performs only slightly better than DSA. This also aligns with our observation that the over-aggressive constraint viola-




Figure 5: Ablation study on sparse random DCOPs

tion conditions in GDBA could severely limit its performance on general-valued DCOPs. Evaporation, on the other hand, also plays an important role in controlling the magnitude of the penalty values. Without it, DGLS suffers from unbounded penalty growth and substantially slower convergence compared to the full DGLS. Interestingly, compared to DGLS w/o AVC, DGLS w/o evaporation exhibits better performance, which also highlights the advantage of our adaptive violation over the fixed violation rules in GDBA. Nonetheless, the combination of AVC and evaporation works synergistically to enable effective selective penalization, as demonstrated by the strong performance of DGLS w/o CPU and full DGLS. Finally, CPU ensures that agents optimize coherently given the changing cost modifiers, leading to moderate but consistent improvements.

5 Conclusion

As an instantiation of GLS on DCOPs, GDBA aims to help local search break out of local optima by increasing the penalty values of violated constraints. However, its empirical benefits remain marginal on general-valued problems. This suboptimal performance stems from three key issues: over-aggressive constraint violation conditions, unbounded penalty accumulation, and uncoordinated penalty updates. Such pathologies lead to ubiquitous over-penalization and indiscriminate penalization, ultimately undermining the intended benefits of breakout.

We therefore present DGLS, a novel GLS framework for DCOPs that effectively addresses these issues by incorporating an adaptive violation condition to selectively penalize constraints with high cost, a penalty evaporation mechanism to control the magnitude of penalization, and a synchronization scheme for coordinated penalty updates. Theoretically, we show that the penalty values of our DGLS are bounded, and agents play a potential game where the potential function is the total augmented cost given the current cost modifiers. Our extensive empirical evaluations on various standard benchmarks confirm the superiority of DGLS over the existing local search heuristics, as well as the state-of-the-art Damped Max-sum on both general-valued and cost-structured problems.

Acknowledgments

This research is supported by the Ministry of Education, Singapore, under its MOE AcRF Tier 2 Award MOE-T2EP20223-0003. Xinrun Wang is supported by Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 1 grant (No. MSS24C005). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of the Ministry of Education, Singapore.

References

- Barabási, A.-L.; and Albert, R. 1999. Emergence of scaling in random networks. *Science*, 286(5439): 509–512.
- Cai, S.; and Lei, Z. 2020. Old techniques in new ways: Clause weighting, unit propagation and hybridization for maximum satisfiability. *Artificial Intelligence*, 287: 103354.
- Chapman, A. C.; Rogers, A.; Jennings, N. R.; and Leslie, D. S. 2011. A unifying framework for iterative approximate best-response algorithms for distributed constraint optimization problems. *The Knowledge Engineering Review*, 26(4): 411–444.
- Chechetka, A.; and Sycara, K. 2006. No-commitment branch and bound search for distributed constraint optimization. In *AAMAS*, 1427–1429.
- Chen, Z.; Deng, Y.; Wu, T.; and He, Z. 2018. A class of iterative refined Max-sum algorithms via non-consecutive value propagation strategies. *Autonomous Agents and Multi-Agent Systems*, 32(6): 822–860.
- Chen, Z.; Zhang, W.; Deng, Y.; Chen, D.; and Li, Q. 2020. RMB-DPOP: Refining MB-DPOP by reducing redundant inference. In *AAMAS*, 249–257.
- Cohen, L.; Galiki, R.; and Zivan, R. 2020. Governing convergence of Max-sum on DCOPs through damping and splitting. *Artificial Intelligence*, 279: 103212.
- Dechter, R. 1998. Bucket elimination: A unifying framework for probabilistic inference. In *Learning in Graphical Models*, volume 89 of *NATO ASI Series*, 75–104. Springer.
- Deng, Y.; Chen, Z.; Chen, D.; Jiang, X.; and Li, Q. 2019. PT-ISABB: A hybrid tree-based complete algorithm to solve asymmetric distributed constraint optimization problems. In *AAMAS*, 1506–1514.
- Farinelli, A.; Rogers, A.; Petcu, A.; and Jennings, N. R. 2008. Decentralised coordination of low-power embedded devices using the Max-sum algorithm. In *AAMAS*, 639–646.
- Fioretto, F.; Pontelli, E.; and Yeoh, W. 2018. Distributed constraint optimization problems and applications: A survey. *Journal of Artificial Intelligence Research*, 61: 623–698.
- Fioretto, F.; Yeoh, W.; Pontelli, E.; Ma, Y.; and Ranade, S. J. 2017. A distributed constraint optimization (DCOP) approach to the economic dispatch with demand response. In *AAMAS*, 999–1007.
- Grinshpoun, T.; Grubshtein, A.; Zivan, R.; Netzer, A.; and Meisels, A. 2013. Asymmetric distributed constraint optimization problems. *Journal of Artificial Intelligence Research*, 47: 613–647.
- Hirayama, K.; and Yokoo, M. 2005. The distributed breakout algorithms. *Artificial Intelligence*, 161(1-2): 89–115.
- Hoang, K. D.; Fioretto, F.; Yeoh, W.; Pontelli, E.; and Zivan, R. 2018. A large neighboring search schema for multi-agent optimization. In *CP*, 688–706.
- Litov, O.; and Meisels, A. 2017. Forward bounding on pseudo-trees for DCOPs and ADCOPs. *Artificial Intelligence*, 252: 83–99.
- Maheswaran, R. T.; Pearce, J. P.; and Tambe, M. 2004. Distributed algorithms for DCOP: A graphical-game-based approach. In *ISCA PDCS*, 432–439.
- Maheswaran, R. T.; Tambe, M.; Bowring, E.; Pearce, J. P.; and Varakantham, P. 2004. Taking DCOP to the Real World: Efficient Complete Solutions for Distributed Multi-Event Scheduling. In *AAMAS*, 310–317.
- Modi, P. J.; Shen, W.-M.; Tambe, M.; and Yokoo, M. 2005. Adopt: Asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence*, 161(1-2): 149–180.
- Monteiro, T. L.; Pujolle, G.; Pellenz, M. E.; Penna, M. C.; and Souza, R. D. 2012. A multi-agent approach to optimal channel assignment in WLANs. In *WCNC*, 2637–2642.
- Netzer, A.; Grubshtein, A.; and Meisels, A. 2012. Concurrent forward bounding for distributed constraint optimization problems. *Artificial Intelligence*, 193: 186–216.
- Nguyen, D. T.; Yeoh, W.; Lau, H. C.; and Zivan, R. 2019. Distributed Gibbs: A linear-space sampling-based DCOP algorithm. *Journal of Artificial Intelligence Research*, 64: 705–748.
- Okamoto, S.; Zivan, R.; and Nahon, A. 2016. Distributed breakout: Beyond satisfaction. In *IJCAI*, 447–453.
- Ottens, B.; Dimitrakakis, C.; and Faltings, B. 2017. DUCT: An upper confidence bound approach to distributed constraint optimization problems. *ACM Transactions on Intelligent Systems and Technology*, 8(5): 69:1–69:27.
- Pearce, J. P.; and Tambe, M. 2007. Quality guarantees on k-optimal solutions for distributed constraint optimization problems. In *IJCAI*, 1446–1451.
- Pertzovsky, A.; Zivan, R.; and Agmon, N. 2024. Collision avoiding max-sum for mobile sensor teams. *Journal of Artificial Intelligence Research*, 79: 1281–1311.
- Petcu, A.; and Faltings, B. 2005. A scalable method for multiagent constraint optimization. In *IJCAI*, 266–271.
- Petcu, A.; and Faltings, B. 2007. MB-DPOP: A new memory-bounded algorithm for distributed optimization. In *IJCAI*, 1452–1457.
- Voudouris, C.; Tsang, E. P.; and Alsheddy, A. 2010. Guided local search. In *Handbook of metaheuristics*, 321–361. Springer.
- Yeoh, W.; Felner, A.; and Koenig, S. 2010. BnB-ADOPT: An asynchronous branch-and-bound DCOP algorithm. *Journal of Artificial Intelligence Research*, 38: 85–133.

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

Zhang, W.; Wang, G.; Xing, Z.; and Wittenburg, L. 2005. Distributed stochastic search and distributed breakout: Properties, comparison and applications to constraint optimization problems in sensor networks. *Artificial Intelligence*, 161(1-2): 55–87.

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

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