

Improved Bounded Max-Sum for Distributed Constraint Optimization

Emma Rollon and Javier Larrosa¹

Abstract. Bounded Max-Sum is a message-passing algorithm for solving Distributed Constraint Optimization Problems able to compute solutions with a guaranteed approximation ratio. Although its approximate solutions were empirically proved to be within a small percentage of the optimal solution on low and moderately dense problems, in this paper we show that its theoretical approximation ratio is overestimated, thus overshadowing its good performance. We propose a new algorithm, called Improved Bounded Max-Sum, whose approximate solutions are at least as good as the ones found by Bounded Max-Sum and with a tighter approximation ratio. Our empirical evaluation shows that the new approximation ratio is significantly tighter.

1 Introduction

Decentralised coordination techniques are a very important topic of research. A common approach is to cast the problem as a *multi-agent distributed constraint optimization problem* (DCOP), where the possible actions that agents can take are associated with *variables* and the utility for taking joint actions are encoded with (soft) *constraints* [10]. The set of constraints define a global utility function $F(x)$ to be optimized via decentralised coordination of the agents. In general, complete algorithms [7, 6, 9] (i.e. algorithms that find the true optimum) exhibit an exponentially increasing coordination overhead, which makes them useless in many practical situations.

Approximate algorithms constitute a very interesting alternative. They require little computation and communication at the cost of sacrificing optimality. There are several examples showing that they can provide solutions which are very close to optimality [3, 5]. However, this observation can only be verified on small toy instances, because it requires the computation of the true optimal to compare with, and it is not available in real-size real-world situations.

A significant breakthrough along this line of work was the Bounded Max-Sum algorithm (BMS) [10]. This algorithm comes with a guarantee approximation ratio $\tilde{\rho}$, meaning that its approximate solution \tilde{x} has a utility $F(\tilde{x})$ which is no more than a factor $\tilde{\rho} \geq 1$ away from the optimum (i.e. $F(\tilde{x}) \leq F(x^*) \leq \tilde{\rho}F(\tilde{x})$). Clearly, large values of $\tilde{\rho}$ reflect lack of confidence in the solution \tilde{x} . There are two possible reasons for a large $\tilde{\rho}$: *i*) the algorithm failed in finding a solution close to the optimal, *ii*) the approximation ratio is not tight. Clearly, if we want $\tilde{\rho}$ to be our measure of confidence about the quality of \tilde{x} , we want a tight $\tilde{\rho}$ (i.e. $F(x^*) \approx \tilde{\rho}F(\tilde{x})$). Thus, the quality of the approximation ratio is a matter of the utmost importance.

In this paper we propose an improvement of BMS with approximation ratio ρ . We theoretically show that it is always better than the previous one (i.e., $\rho \leq \tilde{\rho}$). Moreover, our experiments show that, in practice, ρ is much tighter than $\tilde{\rho}$.

2 Preliminaries

In this Section we review the main elements to contextualize our work. Definitions and notation is borrowed almost directly from [10]. We urge the reader to visit that reference for more details and examples.

2.1 DCOP

A *Distributed Constraint Optimization Problem* (DCOP) is a quadruple $P = (\mathbf{A}, \mathbf{X}, \mathbf{D}, \mathbf{F})$, where $\mathbf{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_r\}$ is a set of agents, and $\mathbf{X} = \{x_1, \dots, x_n\}$ and $\mathbf{D} = \{\mathbf{d}_1, \dots, \mathbf{d}_n\}$ are variables and domains. $\mathbf{F} = \{f_1, \dots, f_e\}$ is a set of cost functions. The objective function is,

$$F(x) = \sum_{j=1}^e f_j(x^j)$$

where $x^j \subseteq \mathbf{X}$ is the scope of f_j . A *solution* is a complete assignment \mathbf{x} . An *optimal solution* is a complete assignment \mathbf{x}^* such that $\forall \mathbf{x}, F(\mathbf{x}^*) \geq F(\mathbf{x})$. The usual task of interest is to find \mathbf{x}^* through the coordination of the agents.

In the applications under consideration, the agents search for the optimum via decentralised coordination. We assume that each agent can control only its local variable(s) and has knowledge of, and can directly communicate with, a few neighboring agents. Two agents are neighbors if there is a relationship connecting variables and functions that the agents control.

The structure of a DCOP problem $P = (\mathbf{A}, \mathbf{X}, \mathbf{D}, \mathbf{F})$ is represented by its associated factor graph. A *factor graph* is a bipartite graph having a variable node for each variable $x_i \in \mathbf{X}$, a factor node for each local function $f_j \in \mathbf{F}$, and an edge connecting variable node x_i to factor node f_j if and only if x_i is an argument of f_j .

2.2 Max-Sum Algorithm

The *Max-Sum* algorithm [2, 1] is a message-passing algorithm for solving DCOP problems. It operates over a factor graph by sending functions (a.k.a., messages) along its edges. Edge (i, j) has associated two messages $q_{i \rightarrow j}$, from variable node x_i to function node f_j , and $r_{j \rightarrow i}$, from function node f_j to variable node x_i . These messages are defined as follows:

¹ Departament de Llenguatges i Sistemes Informàtics, Universitat Politècnica de Catalunya, Spain.

- **From variable to function:**

$$q_{i \rightarrow j}(x_i) = \alpha_{ij} + \sum_{k \in \mathcal{M}_i \setminus j} r_{k \rightarrow i}(x_i)$$

where \mathcal{M}_i is a vector of function indexes, indicating which function nodes are connected to variable node x_i , and α_{ij} is a normalizing constant to prevent the messages from increasing endlessly in cyclic graphs.

- **From function to variable:**

$$r_{j \rightarrow i}(x_i) = \max_{x^j \setminus x_i} \{f_j(x^j) + \sum_{k \in \mathcal{N}_j \setminus i} q_{k \rightarrow i}(x_i)\}$$

where \mathcal{N}_j is a vector of variable indexes, indicating which variable nodes are connected to function node f_j and $x^j \setminus x_i = \{x_k \mid k \in \mathcal{N}_j \setminus i\}$

Max-Sum is a distributed synchronous algorithm, since the agent controlling node i has to wait to receive messages from all its neighbors but j , to be able to compute (and send) its message to j . When the factor graph is cycle free, the algorithm is guaranteed to converge to the global optimal solution. Once the convergence is reached, each variable node can compute function,

$$z_i(x_i) = \max_{x_i} \sum_{k \in \mathcal{M}_i} r_{k \rightarrow i}(x_i)$$

The optimal solution is $\max_{x_i} \{z_i(x_i)\}$ and the optimal assignment $\mathbf{x}_i^* = \arg \max_{x_i} \{z_i(x_i)\}$. When the factor graph is cyclic, the algorithm may not converge to the optimum and only provides an approximation.

3 Bounded Max-Sum Algorithm

The *Bounded Max-Sum* algorithm (BMS) [10], is an approximation algorithm built on the Max-Sum algorithm. From a possibly cyclic problem P , the idea is to remove cycles in its factor graph by ignoring dependencies between functions and variables which have the least impact on the solution quality, producing a new acyclic problem \tilde{P} . Then, Max-Sum is used to optimally solve \tilde{P} while simultaneously computing the approximation ratio $\tilde{\rho}$. A more detailed description follows. For the sake of simplicity, we will restrict ourselves to the case of binary functions $f_j(x_i, x_k)$. The extension to general functions is direct. The algorithm works in three phases, each one implementable in a decentralised manner (see [10] for further details):

- **Relaxation Phase:** First, the algorithm weights each edge (i, j) of the original factor graph as,

$$w_{ij} = \max_{x_k} \{ \max_{x_i} f_j(x_i, x_k) - \min_{x_i} f_j(x_i, x_k) \}$$

Then, it finds a maximum spanning tree T . Let W be the sum of weights of the removed edges (i.e. $W = \sum_{(i,j) \notin T} w_{ij}$). Next, the original problem P is transformed into an acyclic one \tilde{P} having the spanning tree T as factor graph. This is done as follows: for each edge (i, j) in the original graph that does not belong to the tree, the cost function $f_j(x_i, x_k)$ is transformed into another function $\tilde{f}_j(x_k)$ defined as,

$$\tilde{f}_j(x_k) = \min_{x_i} f_j(x_i, x_k)$$

Note that the objective function of \tilde{P} is

$$\tilde{F}(x) = \sum_{(i,j),(k,j) \in T} f_j(x_i, x_k) + \sum_{(i,j) \notin T} \tilde{f}_j(x_k)$$

- **Solving Phase:** BMS solves \tilde{P} with Max-Sum. Let $\tilde{\mathbf{x}}$ be the solution of this problem. Since the factor graph of \tilde{P} is acyclic, $\tilde{\mathbf{x}}$ is its optimal assignment.
- **Bounding Phase:** In [10], it is proved that,

$$F(\tilde{\mathbf{x}}) \leq F(\mathbf{x}^*) \leq \tilde{F}(\tilde{\mathbf{x}}) + W$$

We can rewrite the previous upper bound expression as,

$$F(\mathbf{x}^*) \leq \frac{\tilde{F}(\tilde{\mathbf{x}}) + W}{F(\tilde{\mathbf{x}})} F(\tilde{\mathbf{x}})$$

Therefore, the algorithm computes $\tilde{\rho} = \frac{\tilde{F}(\tilde{\mathbf{x}}) + W}{F(\tilde{\mathbf{x}})}$, which is a guarantee approximation ratio.

4 Improved BMS

4.1 Theoretical elements

Consider an edge (i, j) in the original factor graph that does not belong to the spanning tree. We define $\hat{f}_j(x_k)$ as,

$$\hat{f}_j(x_k) = \max_{x_i} f_j(x_i, x_k)$$

Let \hat{P} denote the problem containing the not modified functions $f_j(x_i, x_k)$ (for $(i, j), (k, j) \in T$) and the $\hat{f}_j(x_k)$ functions (for $(i, j) \notin T$). Note that \hat{P} and \tilde{P} have the same acyclic factor graph. Note as well that the objective function of \hat{P} is

$$\hat{F}(x) = \sum_{(i,j),(k,j) \in T} f_j(x_i, x_k) + \sum_{(i,j) \notin T} \hat{f}_j(x_k)$$

We can solve \hat{P} with Max-Sum. Let $\hat{\mathbf{x}}$ be the optimal solution of this problem. It is obvious that $F(\hat{\mathbf{x}})$ is a lower bound of $F(\mathbf{x}^*)$. Furthermore, as we prove next, $\hat{F}(\hat{\mathbf{x}})$ is an upper bound of $F(\mathbf{x}^*)$. Therefore, $\hat{\rho} = \frac{\hat{F}(\hat{\mathbf{x}})}{F(\hat{\mathbf{x}})}$ is a guarantee approximation ratio.

Theorem 1 $F(\mathbf{x}^*) \leq \hat{F}(\hat{\mathbf{x}})$.

Proof By definition, $F(\mathbf{x}^*) = \sum_{(i,j),(k,j) \in T} f_j(\mathbf{x}_i^*, \mathbf{x}_k^*) + \sum_{(i,j) \notin T} f_j(\mathbf{x}_i^*, \mathbf{x}_k^*)$. Since for all f_j we have that $f_j(x_i, x_k) \leq \max_{x_i} f_j(x_i, x_k)$, then

$$F(\mathbf{x}^*) \leq \sum_{(i,j),(k,j) \in T} f_j(\mathbf{x}_i^*, \mathbf{x}_k^*) + \sum_{(i,j) \notin T} \max_{x_i} f_j(x_i, \mathbf{x}_k^*) = \hat{F}(\mathbf{x}^*)$$

From the optimality of $\hat{\mathbf{x}}$, we know that $\hat{F}(\mathbf{x}^*) \leq \hat{F}(\hat{\mathbf{x}})$, which proves the theorem.

Next, we show that $\hat{F}(\hat{\mathbf{x}})$ is a tighter upper bound than $\tilde{F}(\tilde{\mathbf{x}}) + W$.

Theorem 2 $\hat{F}(\hat{\mathbf{x}}) \leq \tilde{F}(\tilde{\mathbf{x}}) + W$.

Proof The proof is direct once it has been noted that for all $f_j(x_i, x_k)$,


$$\hat{f}_j(x_k) \leq \tilde{f}_j(x_k) + w_{ij}$$

which we prove next. By definition, the previous equation corresponds to,

$$\max_{x_i} f_j(x_i, x_k) \leq \min_{x_i} f_j(x_i, x_k) + \max_{x_k} \{ \max_{x_i} f_j(x_i, x_k) - \min_{x_i} f_j(x_i, x_k) \}$$

which can be rewritten as,

$$\max_{x_i} f_j(x_i, x_k) - \min_{x_i} f_j(x_i, x_k) \leq \max_{x_k} \{ \max_{x_i} f_j(x_i, x_k) - \min_{x_i} f_j(x_i, x_k) \}$$



x_1	x_2	f_1
a	a	10
a	b	10
b	a	0
b	b	10

x_1	x_2	f_2
a	a	10
a	b	0
b	a	0
b	b	10

x_1	\tilde{f}_1
a	10
b	0

x_1	\hat{f}_1
a	10
b	10

Figure 1. Example of a factor graph containing cycles and a spanning tree formed by removing the edge between variable node x_2 and function node f_1 .

which clearly holds.

We cannot establish any dominance relation between $\tilde{\rho}$ and $\hat{\rho}$ because there is no dominance between $F(\tilde{\mathbf{x}})$ and $F(\hat{\mathbf{x}})$. However, one way to circumvent this situation is to take $\rho = \frac{\hat{F}(\tilde{\mathbf{x}})}{\max\{F(\tilde{\mathbf{x}}), F(\hat{\mathbf{x}})\}}$. The new ratio ρ dominates $\tilde{\rho}$.

Theorem 3 $\rho \leq \tilde{\rho}$.

Proof Direct from Theorem 2 and the fact that $\max\{F(\tilde{\mathbf{x}}), F(\hat{\mathbf{x}})\} \geq F(\tilde{\mathbf{x}})$.

4.2 IBMS

Improved BMS (IMBS) works, as BMS, in three phases:

- **Relaxation Phase:** IBMS computes the spanning tree T and the relaxed problem \tilde{P} exactly as BMS does. Additionally, IBMS computes the relaxed problem \hat{P} .
- **Solving Phase:** IBMS solves \tilde{P} and \hat{P} with Max-Sum. Let $\tilde{\mathbf{x}}$ and $\hat{\mathbf{x}}$ be the solutions of these problems. The agents will act according to the best solution ($\max\{F(\tilde{\mathbf{x}}), F(\hat{\mathbf{x}})\}$).
- **Bounding Phase:** IBMS computes the approximation ratio $\rho = \frac{\hat{F}(\tilde{\mathbf{x}})}{\max\{F(\tilde{\mathbf{x}}), F(\hat{\mathbf{x}})\}}$.

The computation, storage and communication effort of IBMS is essentially twice that of BMS, because it requires solving two relaxed problems with Max-Sum. Given the low cost of BMS, doubling it seems acceptable. However, when it is not the case, one can always run a weaker version of IBMS ignoring \hat{P} . This weaker version will be exactly as costly as BMS. Its disadvantage is that $\hat{\mathbf{x}}$ is not guaranteed to be better than $\tilde{\mathbf{x}}$. In fact, our experiments show that there is no clear winner among them. Interestingly, the approximation ratio of the weaker version $\hat{\rho}$ is systematically better than the approximation ratio of BMS $\tilde{\rho}$.

Example 1 Consider the problem P given in Figure 1 with two variables $\{x_1, x_2\}$ and two functions $\{f_1, f_2\}$. The spanning tree of its factor graph is given with solid lines (i.e., edge (x_2, f_1) has been removed, shown as a dashed line). Thus, $W = 10$. Functions \tilde{f}_1 and \hat{f}_1 in \tilde{P} and \hat{P} , respectively, are given in the figure. Max-Sum finds assignments $\tilde{\mathbf{x}} = \hat{\mathbf{x}} = (x_1 = a, x_2 = a)$, with utility $\tilde{F}(\tilde{\mathbf{x}}) = \hat{F}(\hat{\mathbf{x}}) = 20$. Their evaluation on the original problem P is $F(\tilde{\mathbf{x}}) = F(\hat{\mathbf{x}}) = 20$. The approximation ratios are $\tilde{\rho} = 1.5$, $\hat{\rho} = 1$, and $\rho = 1$.

5 Empirical Evaluation

The purpose of the experiments is to evaluate the improvement of our upper bound $\hat{F}(\tilde{\mathbf{x}})$ and approximation ratios ρ and $\hat{\rho}$ over the BMS upper bound $\tilde{F}(\tilde{\mathbf{x}}) + W$ and approximation ratio $\tilde{\rho}$, respectively. We

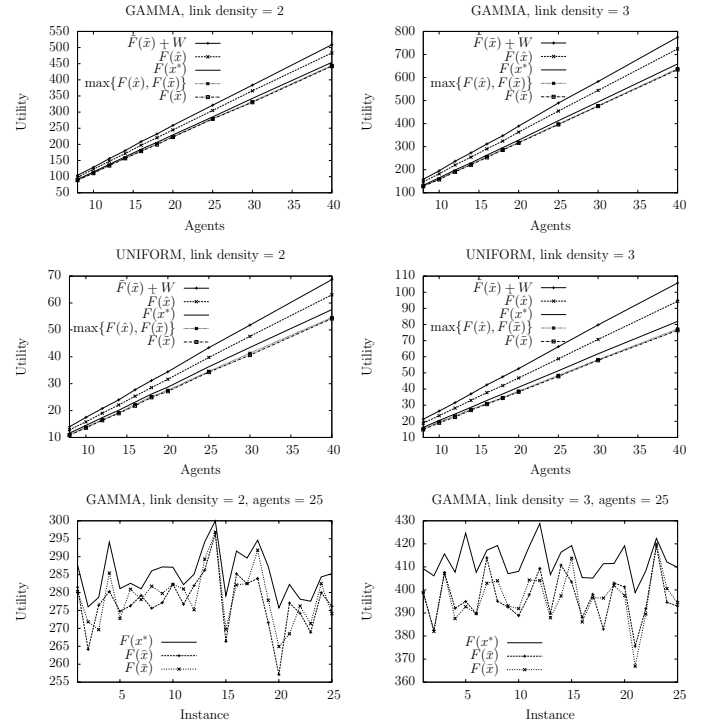


Figure 2. First and second row, bounds obtained by algorithms IBMS and BMS varying the number of agents; third row, lower bound detail for instances with 25 agents and gamma distribution.

consider the same set of problems from the ADOPT repository² used in [10]. These problems represent graph colouring problems with two different link densities (i.e., the average connection per agent) and different number of nodes. Each agent controls one node (i.e., variable), with domain $|\mathbf{d}_i| = 3$, and each edge of the graph represents a pairwise constraint between two agents. Each edge is associated with a random payoff matrix, specifying the payoff that both agents will obtain for every possible combination of their variables' assignments. Each entry of the payoff matrix is a real number sampled from two different distributions: a gamma distribution with $\alpha = 9$ and $\beta = 2$, and a uniform distribution with range $(0, 1)$. For each configuration, we report average values over 25 repetitions. For the sake of comparison, we compute the optimal utility by a complete centralized algorithm, although this value can only be computed up to 12 agents by a complete decentralized algorithm, as shown in [10].

Figure 2 (first and second rows) shows the upper and lower bound

² <http://teamcore.usc.edu/dcop>

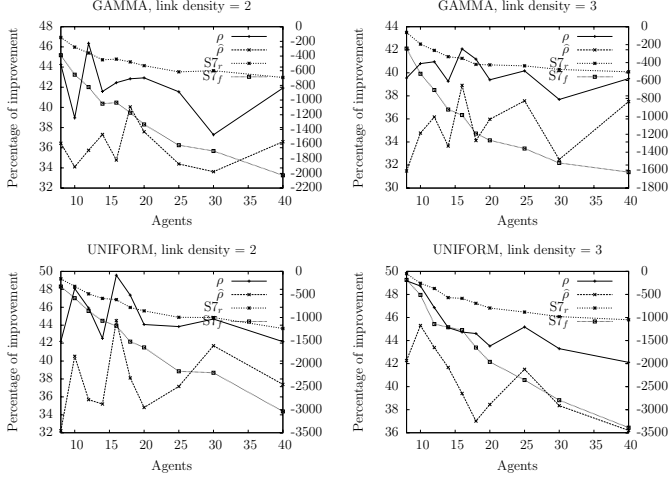


Figure 3. Percentage of improvement of the approximation ratio of IBMS ρ and weaker version of IBMS $\hat{\rho}$ (left y-axis), and percentage of decrease of the 7-size-bounded-distance criteria using the the minimum maximum reward bound ($S7_r$) and the minimum fraction bound ($S7_f$) (right y-axis) over the approximation ratio of BMS $\tilde{\rho}$.

obtained by IBMS (i.e., $\hat{F}(\hat{\mathbf{x}})$ and $\max\{F(\hat{\mathbf{x}}), F(\tilde{\mathbf{x}})\}$, respectively) and BMS (i.e., $\tilde{F}(\tilde{\mathbf{x}})$ and $F(\tilde{\mathbf{x}})$, respectively), along with the optimal utility (i.e., $F(\mathbf{x}^*)$), for the different link densities and payoff distributions. The behavior of both algorithms is very similar across all link densities and payoff distributions. IBMS always computes an upper bound tighter than the one computed by BMS. The improvement is slightly better for the uniform distribution. The lower bounds computed by both algorithms are very close, although IBMS lower bound is slightly better.

Figure 2 (bottom row) shows a detail on the lower bounds $F(\hat{\mathbf{x}})$ and $F(\tilde{\mathbf{x}})$ obtained on each instance of a given parameter configuration. Since the behavior across all number of agents, link densities and payoff distributions is very similar, we only report results on instances with 25 agents and gamma distribution. Both lower bounds are very close, and none of them is consistently better than the other.

Figure 3 shows the percentage of improvement of the approximation ratio of IBMS ρ and the weaker version of IBMS $\hat{\rho}$ over the approximation ratio of BMS $\tilde{\rho}$ (left y-axis). The figure also reports the percentage of deterioration of the approximation ratio of the 7-size-bounded-distance criteria introduced in [11] according to the minimum maximum reward bound ($S7_r$) and the minimum fraction bound ($S7_f$) presented in [12] over the approximation ratio of BMS $\tilde{\rho}$ (right y-axis). Since the relation between the optimal solution of the problem $F(\mathbf{x}^*)$ and an approximation ratio ρ of a given solution \mathbf{x} is $1 \leq \frac{F(\mathbf{x}^*)}{F(\mathbf{x})} \leq \rho$, we compute the improvement of an approximation ratio ρ over $\tilde{\rho}$ as,

$$\frac{(\tilde{\rho} - 1) - (\rho - 1)}{\tilde{\rho} - 1} * 100$$

The improvement of ρ is always higher than 37%, and up to almost 50%. Its mean improvement for the gamma and uniform distributions is higher than 40% and 45%, respectively. The improvement of $\hat{\rho}$ is always higher than 32%, and up to almost 46%. Its mean improvement for the gamma and uniform distributions is higher than 35% and 37%, respectively. Therefore, both IBMS and its weaker version

always significantly outperforms BMS. Recall that the weaker version of IBMS has the same communication demands as BMS. Both approximation ratios $S7_r$ and $S7_f$ are worse than the approximation ratio of BMS $\tilde{\rho}$ (the percentage is always negative). Their quality decreases as the number of agents increases for both distributions.

6 Related Work

There are other two incomplete algorithms that can provide guarantees on the worst-case solution quality of their solutions at design time: *k-optimality* [8] and *t-optimality* [4]. The idea of these algorithms is to form coalitions of agents and to find the local optima solutions for all agents within the coalitions. This local optima is guaranteed to be within a predefined distance from the global optimal solution. Very recently, [11] proposed a framework where different coalition-based local optimality schemes can be described and defined a new criteria called *s-size bounded optimality*. The complexity of these algorithms depend on the number of coalitions and their size. Therefore, in practice, these algorithms are used with relatively small values of their control parameter.

In [10], it was shown that *k-optimality* provided significantly worst quality guarantees than BMS for different values of *k*. As stated in the following proposition, the quality guarantee provided by the 2-size-bounded optimality for binary DCOPs is always higher than 2.

Proposition 1 Let $P = (\mathbf{A}, \mathbf{X}, \mathbf{D}, \mathbf{F})$ be a binary DCOP (i.e., the arity of the functions is at most 2) such that $|\mathbf{X}| > 2$. The quality guarantee ρ satisfies:

$$2 \leq \rho \leq |\mathbf{F}|$$

for all its 2-size-bounded optimal assignments.

Proof According to [11],

$$\rho = \frac{|\mathcal{C}| - nc_*}{cc_*}$$

where:

- \mathcal{C} is a multi-set of subsets of \mathbf{X} , where $C \in \mathcal{C}$ is a coalition;
- $cc_* = \min_{f \in \mathbf{F}} \{nc(f, \mathcal{C})\}$, where $cc(f, \mathcal{C}) = |\{C \in \mathcal{C} \mid var(f) \subseteq C\}|$;
- $nc_* = \min_{f \in \mathbf{F}} \{cc(f, \mathcal{C})\}$, where $nc(f, \mathcal{C}) = |\{C \in \mathcal{C} \mid var(f) \cap C = \emptyset\}|$;

Let f_{ij} be a function with scope $\{x_i, x_j\}$. In a binary DCOP, its 2-size-bounded region is $\mathcal{C} = \{var(f) \mid f \in \mathbf{F}\}$ (i.e., $|\mathcal{C}| = |\mathbf{F}|$), $cc(f, \mathcal{C}) = 1$ for all $f \in \mathbf{F}$, and $nc(f_{ij}, \mathcal{C}) = |\mathbf{F}| - |\{f_{ik} \in \mathbf{F} \mid j \neq k\}| - |\{f_{lj} \in \mathbf{F} \mid l \neq i\}| + 1$ for all $f_{ij} \in \mathbf{F}$. The minimum nc_* is when the DCOP constraint graph is a star because $nc_* = 0$, so that $\rho = |\mathcal{C}|$ is maximum. Note that on a star with $|\mathbf{X}| = 3$, $\rho = |\mathcal{C}| = 2$. The maximum nc_* is when the DCOP constraint graph is a chain with $|\mathbf{X}| > 3$ (note that a chain with $|\mathbf{X}| = 3$ is a star), because $nc_* = |\mathbf{F}| - 3$ so that its $\rho = 3$. Note that a DCOP with two variables is trivially solved to optimality by the 2-bounded size optimality scheme.

Note that for all instances in the empirical evaluation, the quality guarantees of both IBMS and BMS were smaller than 1.4. For those instances, the s-size-bounded-distance provided worse approximation ratios than IBMS and BMS even using the improved minimum maximum reward and minimum fraction bounds and a relatively high value of the parameter *s*. We leave as future work a more extensive empirical comparison of the different algorithms.

7 Conclusions

In this paper we introduced a new algorithm, called Improved Bounded Max-Sum (IBMS), based on the Bounded Max-Sum algorithm. We theoretically proved that its approximation ratio is always better than the previous one, at the only cost of doubling the communication requirements. We also introduced a weaker version of IBMS having the same communication demands as Bounded Max-Sum. Our experiments show that the approximation ratio of both algorithms is significantly tighter.

ACKNOWLEDGEMENTS

The authors are thankful to the authors of [10, 11] for providing us with their implementations, and to the authors of [10] for sharing their benchmarks with us. This work was supported by project TIN2009-13591-C02-01.

REFERENCES

- [1] S. M. Aji and R. J. McEliece, 'The generalized distributive law', *IEEE Transactions on Information Theory*, **46**(2), 325–343, (2000).
- [2] A. Farinelli, A. Rogers, A. Petcu, and N. R. Jennings, 'Decentralised coordination of low-power embedded devices using the max-sum algorithm', in *AAMAS*, pp. 639–646, (2008).
- [3] S. Fitzpatrick and L. Meetrens, 'Distributed coordination through anarchic optimization', in *Distributed Sensor Networks A multiagent perspective*, pp. 257–293. Kluwer Academic, (2003).
- [4] C. Kiekintveld, Z. Yin, A. Kumar, and M. Tambe, 'Asynchronous algorithms for approximate distributed constraint optimization with quality bounds', in *AAMAS*, pp. 133–140, (2010).
- [5] R. J. Maheswaran, J. Pearce, and M. Tambe, 'A family of graphical-game-based algorithms for distributed constraint optimization problems', in *Coordination of Large-Scale Multiagent Systems*, pp. 127–146. Springer-Verlag, (2005).
- [6] R. Mailler and V. R. Lesser, 'Solving distributed constraint optimization problems using cooperative mediation', in *AAMAS*, pp. 438–445, (2004).
- [7] P. J. Modi, W. M. Shen, M. Tambe, and M. Yokoo, 'Adopt: asynchronous distributed constraint optimization with quality guarantees', *Artif. Intell.*, **161**(1-2), 149–180, (2005).
- [8] J. P. Pearce and M. Tambe, 'Quality guarantees on k-optimal solutions for distributed constraint optimization problems', in *IJCAI*, pp. 1446–1451, (2007).
- [9] A. Petcu and B. Faltings, 'A scalable method for multiagent constraint optimization', in *IJCAI*, pp. 266–271, (2005).
- [10] A. Rogers, A. Farinelli, R. Stranders, and N. R. Jennings, 'Bounded approximate decentralised coordination via the max-sum algorithm', *Artif. Intell.*, **175**(2), 730–759, (2011).
- [11] M. Vinyals, E. Shieh, J. Cerquides, J. A. Rodriguez-Aguilar, Z. Yin, M. Tambe, and E. Bowring, 'Quality guarantees for region optimal dcopt algorithms', in *AAMAS*, pp. 133–140, (2011).
- [12] M. Vinyals, E. Shieh, J. Cerquides, J. A. Rodriguez-Aguilar, Z. Yin, M. Tambe, and E. Bowring, 'Reward-based region optimal quality guarantees', in *OPTMAS Workshop*, (2011).