# Competitive Belief Propagation to Efficiently Solve Complex Multi-agent Negotiations with Network Structure

Ivan Marsa-Maestre $^{1(\boxtimes)},$ Jose Manuel Gimenez-Guzman<sup>1</sup>, Enrique de la Hoz<sup>1</sup>, and David Orden<sup>2</sup>

<sup>1</sup> Computer Engineering Department, University of Alcala, Alcalá de Henares, Spain {ivan.marsa,josem.gimenez,enrique.delahoz}@uah.es

<sup>2</sup> Department of Physics and Mathematics, University of Alcala, Alcalá de Henares, Spain david.orden@uah.es

Abstract. This paper focuses on enabling the use of negotiation for complex system optimisation, whose main challenge nowadays is scalability. Although multi-agent automated negotiation has been studied for decades, it is still a challenge to handle in a scalable and efficient manner negotiation problems involving many issues with complex interdependencies. This is a clear obstacle for the use of automated negotiation in complex networks. This paper proposes a novel perspective on the negotiation process as a *competitive belief propagation* process, where the whole negotiation is modelled as a factor graph and distributed belief propagation techniques (BP) are used to yield a solution. We show that the model adequately suits both simple and complex negotiation settings in the literature, and we validate its efficiency and scalability in a challenging, network structured, channel negotiation setting.

Keywords: Belief propagation  $\cdot$  Network-structured negotiations Wi-Fi channel negotiation

### 1 Introduction

A wide range of real-world systems can be modelled as dynamic sets of interconnected nodes [14,20]. The adequate management of complex networked systems is becoming critical for industrialized countries, since they keep growing in size and complexity. An important sub-class involves autonomous, self-interested entities (e.g. drivers in a transportation network). The self-interested nature of the entities in the network causes the network to deviate from socially-optimal behaviour. This leads to problems related to unavailability and inefficient use of resources.

Different fields of research are working on these challenges, but, so far, with only mixed success. Optimization techniques are especially suited to

© Springer International Publishing AG 2017

G. Sukthankar and J. A. Rodriguez-Aguilar (Eds.): AAMAS 2017 Visionary Papers, LNAI 10643, pp. 1–16, 2017.

 $https://doi.org/10.1007/978\text{-}3\text{-}319\text{-}71679\text{-}4\_1$ 

address large-scale systems with an underlying network structure, usually with a "divide and conquer" approach [22,24]. However, their performance severely decreases as the complexity of the system increases [21], and with the presence of autonomous entities which deviate from the globally optimal solution, thus harming the social goal. Negotiation techniques are known to be useful to handle self-interested behaviour [1], but scale poorly with problem size and the intricacies of interdependencies [15].

Belief propagation (BP) is a message-passing technique which has been successfully used to solve optimization problems by modelling problem constraints using a graph structure [9]. In our previous work, we used it to improve the scalability of an agent self-preference exploration during an auction-based negotiation process over preference spaces built of hypercube-shaped constraints [18]. In a more recent work [11], the concept was extended to utility hypergraphs, enabling negotiations with more complex shaped constraints. These works contributed to greatly enhance scalability in complex negotiations by modelling utilities as graphs, but did not explore problems which were graph-structured by themselves. In addition, these works used belief propagation to explore the preferences of each agent separately, thus keeping the BP process as a local optimization process, using it to assist in the local search for solutions during the negotiation settings.

Our goal is to provide a novel perspective of the multi-agent negotiation process as a *competitive belief propagation*, which can help to efficiently handle conflicts in network-structured settings. In this paper, we contribute to this goal in the following way:

- We define the negotiation problem as a factor graph  $F_P$  and the negotiation process as a competitive belief propagation over the factor graph, and we show how the model suits well-known negotiation settings in the literature (Sect. 2).
- We apply this model to a challenging, network structured complex negotiation setting: Wi-Fi channel negotiation (Sects. 3.1 and 3.2).
- We propose a distributed, scalable approach to use competitive belief propagation to solve the problem (Sect. 3.3).

To test our hypothesis and evaluate the validity of our contribution, we have conducted a set of experiments in a realistic scenario setting (Sect. 4). Our experiments show that the belief propagation approach outperforms a classical nonlinear negotiation approach in terms of solution efficiency and performance. The last section summarizes our contributions and sheds light on future challenges and lines of research.

### 2 Negotiation as Competitive Belief Propagation

As we said above, our goal is to provide an alternative perspective on automated negotiation as a competitive belief propagation process, so that this process can be mapped naturally to network structured problems. To understand the rationale of this novel perspective, let's see a simple example of the correspondence between the two problems. Consider the classic negotiation game between a buyer agent B and a seller agent S over the price p of an item, and let us assume that they are going to use an alternating offer protocol for the negotiation [7]. Let  $b_B^0$  and  $b_S^0$  the initial agent offers (or *bids*). For a typical bargaining scenario, we will have  $b_B^0 < b_S^0$ , that is, the buyer wants to pay less than the seller wants to get. During the negotiation, we will expect to see a progressive relaxation of the agents' initial positions (i.e.  $b_B^t \ge b_B^{t-1}$  and  $b_S^t \le b_S^{t-1}$ ) until we reach the point of agreement (i.e.  $b_B^t = b_S^t$ ) or the deadline expires. The relaxation speed of the agents' positions will depend on the different agent bidding strategies (e.g. boulware, conceder, tit-for-tat [2]). This process can easily be seen as a *belief competition*. Both agents start out with different beliefs about how much they want to give or receive for the item ( $b_B^0$  and  $b_S^0$ , respectively), and these beliefs get updated through the subsequents iterations of the protocol until an agreement is reached (i.e. the beliefs of both agents match) or the deadline expires. This belief competition perspective can be formalized into a factor graph and a belief propagation process [9], as we will see in the following subsection.

#### 2.1 Negotiation as a Factorized Optimization Problem

Belief propagation (BP) techniques have been shown as successful heuristics for solving *factorized optimization problems*, that is, problems P of the form

$$minimize \sum_{i \in V} \Phi_i(x_i) + \sum_{c \in C} \Psi_c(x_c) \tag{1}$$

where V is a finite set of variables and C is a finite collection of subsets of V representing constraints.  $\Phi_i$  and  $\Psi_i$  are real-valued functions called, respectively, variable functions and factor functions, representing the impact on the objective function of the value of each independent variable, and of combinations of variable values (i.e. interdependencies between variables).

We can easily map to this kind of problem most of the utility or social welfare functions used in automated negotiation, by choosing adequate expressions for the  $\Phi_i$  and  $\Psi_i$  functions. For instance, in the trivial buyer-seller negotiation example above we could choose

$$\Phi_B = \begin{cases} x - R_B & \text{if } x \le R_B \\ \infty & \text{otherwise} \end{cases} \qquad \Phi_S = \begin{cases} R_S - x & \text{if } x \ge R_S \\ \infty & \text{otherwise} \end{cases}$$
$$\Psi(x_B, x_S) = \begin{cases} 0 & \text{if } x_B = x_S \\ \infty & \text{otherwise} \end{cases}$$

where  $R_B$  and  $R_S$  are the reservation values of buyer and seller, respectively. In this trivial case, each  $\Phi_i$  functions account for the preferences of agent *i*, and the  $\Psi$  function represents that a disagreement is the worst possible outcome. Other  $\Psi$  functions could be chosen depending on the desired outcome of the negotiation. For instance, we could choose a  $\Psi$  function representing a Clarke tax  $(\Psi(x_B, x_S) = \frac{x_S - x_B}{2})$ , or a fairness metric such as the one defined in [8].

A similar translation can be done in a linear-additive multi-issue negotiation setting, such as the one described in [6,17]. Here there would be a  $\Phi_{a,i}(x_{a,i})$ function for each agent *a* and issue *i*, corresponding to the valuation function of each agent for each issue.  $\Psi_i$  functions would be defined for each issue depending on the desired outcome of the negotiation. For instance, to introduce the aforementioned fairness measure we could define:

$$\Psi(x_B, i, x_S, i) = \begin{cases} \frac{(\Phi_B(x_B) + \Phi_S(x_S))^2 + (\Phi_B(x_B) - \Phi_S(x_S))^2}{2} & \text{if } x_B = x_S \\ \infty & \text{otherwise} \end{cases}$$

Apart from the  $\Phi$  and  $\Psi$  functions, we can also derive the corresponding factor graph  $F_P$  of the factorized optimization problem P. This is a bipartite graph with a node per  $\Phi$  and  $\Psi$  function, and links between nodes which share variables.

### 2.2 The Negotiation Process as Belief Propagation

Once the negotiation problem has been expressed as a factorized optimization problem  $F_P$ , we can solve it using belief propagation techniques. In particular, we use the min-sum version of BP described in [9], which we reproduce in Algorithm 1 for convenience.

Algorithm 1. min-sum BP
<b>Input</b> : F: bipartite factor graph with edges $(i, f)$ between variable nodes and
factor nodes representing constraints $N$ : number of iterations
$C: \{c_i\}:$ available color set
<b>Output</b> : S: estimated optimal assignment
Initialize $t = 0$
foreach $edge$ $(i, f)$ in $F$ do
initialize $m_{f \to i}^0(z) \forall z \in C$
end
for $t = 1, 2, \dots, N$ do
foreach $edge$ $(i, f)$ in $F$ do
update $m_{i \to f}^t(z) = \Phi_i(z) + \sum_{k \in f_i \setminus f} m_{k \to i}^{t-1}(z)$
update $m_{f \to i}^t(z) = \min_{y \in C^{ f }, y_i = z} \Psi_f(y) + \sum_{j \in f \setminus i} m_{j \to f}^{t-1}(y_j)$
$\mathbf{end}$
t = t + 1
end
Set the belief function as
$b_i^N = \Phi_i(z) + \sum_{k \in f_i} m_{k \to i}^N(z)$ for each variable node i
Estimate the optimal assignment $S$ as
$\hat{s}_i^{N(z)}$ for each variable node <i>i</i>

The min-sum algorithm is a message-passing algorithm which defines a number of messages  $m_{X\to Y}^t$  to be passed among nodes of the factor graph  $F_P$ throughout the different iterations of the algorithm. We turn the algorithm into a multi-agent negotiation protocol by mapping the graph nodes to agents. In our example, we would map to each negotiating agent A and B the nodes corresponding to its own beliefs/preferences about the issue values, and the nodes corresponding to the  $\Psi$  functions (in this case, incentivizing fairness) would be mapped to a mediator agent M. The negotiation would progress via message passing between factor and variable nodes assigned to different agents at each iteration.

# 3 Scaling Up: BP in the Wi-Fi Negotiation Problem

In our previous work, we proposed Wi-Fi channel assignment as a realistic and challenging benchmark for complex automated negotiations [3,4]. In this setting, different Wi-Fi providers, acting as agents, have to collectively decide how to distribute the channels used by their access points (APs) in order to minimize interference between nodes and thus maximize the utility (i.e., network throughput) for their clients, which will be different kinds of wireless devices (WDs). A Wi-Fi negotiation scenario will be characterized by:

- A Wi-Fi association graph G, which is a geometric graph (i.e., nodes have specific positions in space) with two kinds of vertices, representing APs and WDs. Edges in the graph represent the association of a particular WD to an AP. In Fig. 1 we show a graphical representation of a scenario with 26 APs and 400 WDs.
- An interference graph I, which includes the same vertices as G in the same positions in space, but in this case edges represent potential interferences between devices, and edge weights account for the intensity of these interferences. For detailed description of the Wi-Fi interference model in this setting, the reader is advised to check [13]. The corresponding graph for the aforementioned scenario can be seen in Fig. 2.
- A mapping of access points to different providers, which will be the negotiating agents. The goal of each agent will be to minimize the interference suffered by its APs and their associated WDs. In this paper, we will assume there are two providers.

This is a particularly interesting problem, since it belongs to the family of Frequency Assignment Problems (which has been extensively studied from the perspective of discrete optimization) and it is strongly related to the prominent mathematical graph coloring problem [23] and to distributed constraint optimization models [10]. In the following, we will formally describe the negotiation problem, and the translation of the problem to the belief propagation model.

#### 3.1 Negotiation Domain

For the scope of this work, we assume a multiattribute negotiation domain, where a deal or solution to the problem is defined as the set of attributes (*issues*),



Fig. 1. Wi-Fi association graph G for an scenario with 26 APs and 400 wireless devices.



Fig. 2. Wi-Fi interference graph I for an scenario with 26 APs and 400 wireless devices.

and each one of them can be in a certain range. In our case, for a channel assignment problem with  $n_{AP}$  access points, a solution or deal S can be expressed as  $S = \{s_i | i \in 1, ..., n_{AP}\}$ , where  $s_i \in \{1, ..., 11\}$  represents the assignation of a Wi-Fi channel to the *i*-th access point.

As stated above, we assume that there are two network providers or agents (commonly Internet Service Providers, ISPs), thus APs belong to one of the agents. Each provider only has control over the channel assignment for its own access points. According to this situation,  $P = \{p_1, p_2\}$  will be the set of agents that will negotiate the channel assignment. We find adequate to focus in the two-provider case because there are more works in complex bilateral negotiations than for the multilateral case (three or more agents).

Finally, each one of these agents  $p_i$  will compute its utility  $U_{p_i}$  for a certain solution according to the model described in [4,13]. As we showed in our previous

works, the problem settings (high cardinality of the solution space and attribute interdependence) will make the utility functions highly complex, with multiple local optima.

#### 3.2 Channel Negotiation as a Factorized Optimization Problem

Given that  $U_{p_i}$  depends on the sum of the interferences suffered by their APs and their associated WDs, and that those interferences are caused by nearby APs and WDs using the same channel to transmit, intuitively we need to avoid using the same (or similar) channels in nearby devices. More specifically, we have to avoid using "close" channels in devices whose associated vertices in the interference graph I are connected, specially if the weight of the connecting edge is high.

This is quite similar to the *Threshold Coloring Problem* (TSC) [19], which is depicted in Fig. 3. In this problem, we have an undirected graph and a set of available colors (in the example, red, green and blue), with an associated interference matrix, which assigns an interference value for the occurrence of any pair of colors in any edge of the graph. The goal of the TSC problem is to find a coloring which minimizes the maximum interference per node (the optimal solutions for the example problem can be seen shadowed in the figure). Our hypothesis is that, by translating the problem of channel assignment to this problem, we will find suitable solutions in a reasonable time. We will have as the available color set the different Wi-Fi channels, and as the color interference matrix the co-channel interference index [13].



**Fig. 3.** Example of the Threshold Spectrum Coloring problem (TSC). (Color figure online)

The translation from the Wi-Fi channel negotiation problem to the TSC problem is not straightforward. First of all, the variable/issue sets for both problems are different. TSC assumes all colorings are possible, while in Wi-Fi channel negotiation only channels chosen by the APs are negotiated, since all WDs associated to a given AP will use the same channel to communicate. To account for

this, we propose to do a *compactation* of the interference graph I. That is, we will derive a *compact graph* C from the graph I as follows:

- We will have a vertex in the graph C for every AP vertex in the graph I.
- Two vertices in C will be connected if and only if there was an edge in I; (a) between them, (b) between one of the APs and one of the WDs associated to the other AP, or (c) between WDs of the two APs.

Such a compact graph for the example given in Figs. 1 and 2 is shown in Fig. 4. In addition, we will introduce edge weights in graph C according to two different strategies:

- Uniform BP translation (BPu): we will assume that each existing edge between vertices i and j in the compact graph C has an associated weight  $w_{ij} = 1$ . This is the simplest possible translation (all edges equal), which loses most information from the interference graph I, so we expect it to give the less optimal results, but also to be the most efficient in terms of computation time.
- Weighted BP translation (BPw): in this case, we will assign to each edge between vertices i and j in the compact graph C a weight  $w_{ij}$  equal to the number of edges in graph I between them, between one of the APs and one of the clients associated to the other AP, and between WDs of the two APs. This is a reasonable choice, since it will prioritize the edges between APs which have more potential interferences, but it is still a much more efficient choice computationally than evaluating the real interference between APs and their WDs.

The last step for the formalization is to translate the coloring of the compact graph C to a factorized optimization problem. To do this, we use our vertices representing APs as variables (which can take different values depending on



Fig. 4. Compact graph C for an scenario with 26 APs and 400 wireless devices.

which channel the AP uses to transmit) and the edges between pairs of vertices (which represent interferences between APs) as constraints. According to this, we define the corresponding functions as follows:

$$\Phi_i(s_i) = 0$$
$$\Psi_C(s_i, s_j) = w_{ij}c(s_i, s_j), \forall C \equiv (i, j)$$

That is, we use a constant zero value for each variable function, and we use the product of the weight of the corresponding edge in the compact graph and the co-channel interference between the chosen channels (see [13] for details) for each factor function. With this formulation, we try to mitigate the impact of using "close" channels in close APs, which is coherent with the Wi-Fi interference model. It is worth noting that this formulation differs from the TSC problem, given that here we try to minimize the sum of the contributions for all nodes in the graph, while pure TSC aims to minimize the maximum contribution for any single node in the graph. However, as shown in [19], sum minimization is a good heuristic to minimize the maximum in this context, and therefore successful techniques proposed for TSC can be used here as benchmarks.

Finally, we need to build the factor graph  $F_P$  of our problem, which is a bipartite graph with *variable nodes* in one side of the partition and *factor nodes* corresponding to the constraints in the other side of the partition. Links between both partitions occur between a constraint and the variable nodes it refers to. For instance, in the graph example given in Fig. 3, the resulting factor graph  $F_P$  would be as shown in Fig. 5.



**Fig. 5.** Factor graph  $F_P$  (right) for our example TSC problem.

#### 3.3 A Scalable Negotiation Using Belief Propagation

To solve the factorized optimization problem we have proposed for our Wi-Fi negotiation scenario, we would have to apply the min-sum algorithm for BP [9]. The problem with applying directly the min-sum BP algorithm to our problem is that the algorithm only has correctness and convergence guarantees when the solution is unique and the factor graph is a tree. Although solution uniqueness can be achieved with randomized weights as suggested in [9], most of our scenarios do not create tree factor graphs. The usual junction tree technique used in machine learning to address this problem [25] is not applicable here, because it is centralized, which is not scalable (neither desirable) for competitive negotiation scenarios where sharing of private information should be minimized.

Taking this into account, to ensure convergence and correctness of the algorithm, we propose to divide the factor graph into trees using a distributed, gossip-inspired technique [5]. The technique we propose works as follows:

- All AP nodes in the compact graph C are initialized to the *unassigned* state, which means they do not belong to any tree.
- Nodes in *unassigned* state respond to the behaviour:
- Decide with probability p whether to start a new tree (therefore changing their status to assigned) or to wait a random time.
- Upon receiving a message from an *assigned* neighbour (that is, a neighbour already belonging to a tree), switch to assigned status and acknowledge the membership to the tree.
- Nodes in *assigned* state respond to the behaviour:
- Decide with probability p whether to invite a random subset of its (not already-invited) neighbors to its tree or to wait a random time.

This technique asynchronously divides the compact graph C into a set of disjoint trees, from which tree factor graphs can be derived so that BP converges. Of course, when we work with the resulting set of trees, we lose the information about the influencing factors  $\Psi_{ij}$  corresponding to components  $c_i$  and  $c_j$  which are neighbors in the compact graph but have ended up in different trees. To minimize the impact of this simplification, we iteratively introduce this effect in the functions  $\Phi_i$  of the frontier nodes (that is, the nodes in a tree which are neighbors of nodes in other trees). That is, the belief propagation process is repeated several times in an iterative manner, and at each iteration K the frontier nodes are assigned a variable function  $\Phi_i^K(s_i)$  which is computed as follows:

$$\Phi_i^K(s_i) = \sum_{j \in \aleph(i)} \Psi_{ij}(s_i, \hat{s}_j^{K-1})$$

Where  $\aleph(i)$  is the set of neighbors of component  $c_i$  in the compact graph and  $\hat{s}_j^{K-1}$  is the optimal assignment for neighbor  $c_j$  for the previous execution of the BP algorithm.

Computation of the  $\Phi_i^K$  functions is performed at each corresponding provider agent for the AP *i*. Computation of the  $\Psi_{C_{ij}}^K$  functions when APs *i* and *j* belong to different providers is randomly assigned to one of the provider agents to avoid agent manipulation of the belief propagation process.



Fig. 6. Polytechnic school building plan. (Color figure online)

### 4 Experimental Evaluation

To validate our approach and assess the contribution of our mechanisms, we have conducted a set of experiments of the Wi-Fi real-world setting we used in [3], which uses the real layout of the first floor plant of our University (Fig. 6). The real positions of deployed APs are displayed with green dots, ranging their signal coverage from red (high coverage) to light blue (very low coverage). Note that the center of the plan represents a central courtyard, so it has low signal coverage. For the position of WDs we have considered that we have users attending classes in classrooms and also some students are located randomly in the building (resting, in the cafeteria, studying...). For this last group of students, we have considered that there are 100 students randomly located in the building following a uniform distribution. For the students in classrooms, we have tested several scenarios varying randomly the ratio of classrooms being used ( $\rho$ , with  $\rho \in [0.25, 0.5, 0.75, 1.0]$ ). As there are 48 classrooms in the building, we have considered scenarios with 12, 24, 36 and 48 classrooms. For each classroom, we have deployed 25 students in each one randomly using a normal distribution around the center of each classroom and a standard deviation normalized to the size of the scenario of 0.05. In Table 1 we show a summary of the real-world scenarios under study. Finally, as the specific random classrooms under use could affect the results, we have tested three experiments for each value of  $\rho$ , so the total number of deployments studied has been 12. In each setting, APs have been randomly assigned to the two providers.

We are interested in evaluating the performance of BPu and BPw in comparison with the well-known technique called *Simulated Annealing* (SA) [12,16]. For a further description of how this technique has been deployed, see [4]. More specifically, we are interested in evaluating the performance of these techniques in terms of the normalized utility  $(U_n)$  that they can achieve in a certain

Scenario	ρ	# Classrooms	# WD
1, 2, 3	0.25	12	400
4, 5, 6	0.5	12	700
7, 8, 9	0.75	12	1000
10, 11, 12	1.0	48	1300

0.34 0.34 0.32 0.32 0.3 0.3 0.28 0.28 0.26 0.26 5 0.24 5 0.24 0.22 0.22 0.2 0.2 0.18 0.18 • BPu • BPu - BPw 0.16 BPw 0.16 -× SA ×−× SA 0.14 0.14 1 1.2 1.4 1.6 1.8 2.4 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1 1.1 1.2 1.3 1.4 1.5 0.2 0.4 0.6 0.8 2 2.2 2.6 1.6 (a) Scenario 1. (b) Scenario 2. 0.34 0.3 0.32 0.28 0.3 0.26 0.28 0.24 0.26  $\stackrel{\alpha}{\eqsim} 0.24$  $\stackrel{a}{\supset} 0.22$ 0.22 0.2 0.2 0.18 0.18 BPu BPu 0.16 0.16 - BPw BPw ×SA K SA 0.140.14 1.5 0.5 1.5 2.5 2 2.5 2 3.5 0.5 4 4.5 0 (c) Scenario 3. (d) Scenario 4. 0.3 0.3 0.28 0.28 0.26 0.26 0.24 0.24 0.22  $\stackrel{e}{\supset} 0.22$  ${\rm U}_n$ 0.2 0.18 0.2 0.16 0.18 0.14 BPu • BPu 0.16 • BPw • BPw 0.12 → SA ×→× SA 0.1 0.14 -. . . 0.5 1.5 2.5 3 3.5 1.5 2 2.5 3 3.5 4.5 1 4.: 0.5 (e) Scenario 5. (f) Scenario 6.

 Table 1. Summary of scenarios of the real-world setting.

Fig. 7. Normalized utility versus time for the different techniques.



Fig. 8. Normalized utility versus time for the different techniques.

computation time. The different values for the computation time have been obtained running the techniques with a different number of iterations.

Note that the normalized utility is defined as the sum of utilities for all nodes in the network (APs and WDs) divided by the graph order, i.e. divided by the number of nodes in the graph. Figures 7 and 8 show this comparison for the 12 scenarios under study. Results show that, except for the shortest runs where BP obtains worse results than SA, in almost all cases, BP is able to obtain a better performance (higher  $U_n$ ) than SA for the same computation time.

Comparing BPu with BPw we can conclude that BPw always outperforms BPu. This expected result is due to the fact that the compact graph of BPw includes more information than its counterpart of BPu. As a consequence of these results, we consider that the use of Belief Propagation, specially in its BPw setting, is very useful, as its efficiency is higher than the well-known, successful approach SA.

# 5 Conclusions and Future Work

Our research attempts to enable the use of complex automated negotiations in the management of complex systems with network structure, which are of increasing interest in many disciplines. One of the biggest challenges for this is the scalability of complex negotiation mechanisms when facing the large utility spaces and complex interdependencies of such systems. To address this challenge, in this paper we propose a novel perspective for the negotiation process as competitive belief propagation, which maps naturally to settings with a graph structure. We also propose an efficient mechanism to implement this competitive belief propagation process in large settings, which takes advantage of gossip-like techniques. Finally, we validate our approach on a realistic and challenging setting: Wi-Fi channel negotiation. Experiments show that our approach achieves better results that well-known successful nonlinear negotiation techniques in less computation time, which is a significant advance for the success of negotiation mechanisms in these settings.

Although our experiments yield satisfactory results, there is still plenty of work to be done in this area. We are interested in evaluating more sophisticated strategies than BPu and BPw for weighting the compact graph, in order to get as close as possible to the real interference model without imposing too much computational complexity in the mechanism. We want also to study the influence of different graph properties (e.g. diameter) in the performance of the BP techniques. Finally, we are interested in evaluating the strategic properties of the mechanisms, to see how the belief propagation process performs when agents are allowed to "lie" in their messages in order to try to influence the outcome of the mechanism to their advantage.

Acknowledgments. This work has been supported by the Spanish Ministry of Economy and Competitiveness grants TIN2016-80622-P, TIN2014-61627-EXP, MTM2014-54207 and TEC2013-45183-R, and by the University of Alcala through CCG2016/EXP-048.

# References

- An, B., Lesser, V., Sim, K.M.: Strategic agents for multi-resource negotiation. Auton. Agent. Multi-Agent Syst. 23(1), 114–153 (2011)
- Baarslag, T., Dirkzwager, A., Hindriks, K.V., Jonker, C.M.: The significance of bidding, accepting and opponent modeling in automated negotiation. In: Proceedings of the Twenty-First European Conference on Artificial Intelligence, pp. 27–32. IOS Press (2014)

- De La Hoz, E., Gimenez-Guzman, J.M., Marsa-Maestre, I., Orden, D.: A realistic scenario for complex automated nonlinear negotiation: Wi-Fi channel assignment. In: Proceedings of the the Ninth International Workshop on Agent-based Complex Automated Negotiations (ACAN2016), Singapore (2016)
- 4. De La Hoz, E., Marsa-Maestre, I., Gimenez-Guzman, J.M., Orden, D., Klein, M.: Multi-agent nonlinear negotiation for Wi-Fi channel assignment. In: Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2017, International Foundation for Autonomous Agents and Multiagent Systems, Sao Paulo, Brazil (2017)
- Demers, A., Greene, D., Hauser, C., Irish, W., Larson, J., Shenker, S., Sturgis, H., Swinehart, D., Terry, D.: Epidemic algorithms for replicated database maintenance. In: Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing, PODC 1987, New York, NY, USA, pp. 1–12 (1987)
- Faratin, P., Sierra, C., Jennings, N.R.: Using similarity criteria to make issue tradeoffs in automated negotiations. Artif. Intell. 142(2), 205–237 (2002)
- Fatima, S., Kraus, S., Wooldridge, M.: Principles of Automated Negotiation. Cambridge University Press, Cambridge (2014)
- Fujita, K., Ito, T., Klein, M.: A secure and fair protocol that addresses weaknesses of the nash bargaining solution in nonlinear negotiation. Group Decis. Negot. 21(1), 29–47 (2012)
- Gamarnik, D., Shah, D., Wei, Y.: Belief propagation for min-cost network flow: convergence and correctness. Oper. Res. 60(2), 410–428 (2012)
- Grubshtein, A., Meisels, A.: A distributed cooperative approach for optimizing a family of network games. In: Brazier, F.M.T., Nieuwenhuis, K., Pavlin, G., Warnier, M., Badica, C. (eds.) Intelligent Distributed Computing V: Proceedings of the 5th International Symposium on Intelligent Distributed Computing - IDC 2011, pp. 49–62. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-24013-3\_6
- 11. Hadfi, R., Ito, T.: Complex multi-issue negotiation using utility hyper-graphs. JACIII **19**(4), 514–522 (2015)
- Hattori, H., Klein, M., Ito, T.: Using iterative narrowing to enable multi-party negotiations with multiple interdependent issues. In: Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS 2007, pp. 247:1–247:3. ACM, New York (2007)
- de la Hoz, E., Gimenez-Guzman, J.M., Marsa-Maestre, I., Orden, D.: Automated negotiation for resource assignment in wireless surveillance sensor networks. Sensors 15(11), 29547–29568 (2015)
- Kinney, R., Crucitti, P., Albert, R., Latora, V.: Modeling cascading failures in the North American power grid. Eur. Phys. J. B-Condens. Matter Complex Syst. 46(1), 101–107 (2005)
- Klein, M., Faratin, P., Sayama, H., Bar-Yam, Y.: Negotiating complex contracts. Group Decis. Negot. 12(2), 111–125 (2003)
- Lang, F., Fink, A.: Learning from the metaheuristics: protocols for automated negotiations. Group Decis. Negot. 24(2), 299–332 (2015)
- Lopez-Carmona, M.A., Marsa-Maestre, I., Ibañez, G., Carral, J.A., Velasco, J.R.: Improving trade-offs in automated bilateral negotiations for expressive and inexpressive scenarios. J. Intell. Fuzzy Syst. 21(3), 165–174 (2010)
- Marsa-Maestre, I., Lopez-Carmona, M.A., Velasco, J.R., de la Hoz, E.: Effective bidding and deal identification for negotiations in highly nonlinear scenarios. In: Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 2, AAMAS 2009, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, pp. 1057–1064 (2009)

- Orden, D., Marsá-Maestre, I., Giménez-Guzmán, J.M., de la Hoz, E.: Spectrum graph coloring and applications to WiFi channel assignment. CoRR abs/1602.05038 (2016)
- Osorio, C., Bierlaire, M.: Mitigating network congestion: analytical models, optimization methods and their applications. In: 90th Annual Meeting, No. EPFL-TALK-196049 (2011)
- Pelikan, M., Sastry, K., Goldberg, D.E.: Multiobjective estimation of distribution algorithms. In: Pelikan, M., Sastry, K., CantúPaz, E. (eds.) Scalable Optimization via Probabilistic Modeling, pp. 223–248. Springer, Heidelberg (2006). https://doi. org/10.1007/978-3-540-34954-9\_10
- Schaller, B.: New York City's congestion pricing experience and implications for road pricing acceptance in the United States. Transp. Pol. 17(4), 266–273 (2010)
- Tuza, Z., Gutin, G., Plurnmer, M., Tucker, A., Burke, E., Werra, D., Kingston, J.: Colorings and related topics. Handbook of Graph Theory. Discrete Mathematics and Its Applications, pp. 340–483. CRC Press, Boca Raton (2003)
- Vytelingum, P., Ramchurn, S.D., Voice, T.D., Rogers, A., Jennings, N.R.: Trading agents for the smart electricity grid. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1, International Foundation for Autonomous Agents and Multiagent Systems, pp. 897–904 (2010)
- Zheng, L., Mengshoel, O.: Optimizing parallel belief propagation in junction treesusing regression. In: Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2013, ACM, Chicago, Illinois, USA, pp. 757–765 (2013)