Scalable decentralized supply chain formation through binarized belief propagation

(Extended Abstract)

Toni Penya-Alba, Jesus Cerquides, Juan A. Rodriguez-Aguilar
IIIA - CSIC
Campus de la UAB, E-08193 Bellaterra, Spain
{tonipenya, cerquide, jar}@iiia.csic.es

Meritxell Vinyals
Department of Computer Science
University of Verona
Strada le Grazie, 15, Verona, Italy
meritxell.vinalssalgado@univr.it

ABSTRACT

Supply Chain Formation (SCF) is the process of determining the participants in a supply chain, who will exchange what with whom, and the terms of the exchanges. Decentralized SCF appears as a highly intricate task because agents only possess local information, have limited knowledge about the capabilities of other agents, and prefer to preserve privacy. Very recently, the decentralized SCF problem has been cast as an optimization problem that can be efficiently approximated using max-sum loopy belief propagation. Unfortunately, the memory and communication requirements of this approach largely hinder its scalability. This paper presents a novel encoding of the problem into a binary factor graph (containing only binary variables) along with an alternative algorithm. These allow to scale up to form supply chains in markets with higher degrees of competition.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems

General Terms

Algorithms, Economics, Experimentation

Keywords

Supply chain, belief propagation, scalability

1. INTRODUCTION

According to [4], "Supply Chain Formation (SCF) is the process of determining the participants in a supply chain, who will exchange what with whom, and the terms of the exchanges". Although intractable [3], the SCF problem has been widely tackled by the multi-agent systems (MAS) literature, mainly through centralized auction-based approaches [5, 1]. Furthermore, as argued in [4], even when the computation is tractable, no single entity may have global allocative authority to compute allocations over the entire supply chain (SC). To overcome these limitations, a decentralized manner to solve the problem is proposed in [4]. More recently, Winsper et al. [6] cast the decentralized SCF problem into an optimization problem that can be approximated using max-sum loopy belief propagation (LBP)[1]. Unfortunately, the memory and communication requirements of this approach hinder its scalability.

In this paper we propose a novel approach to the decentralized SCF problem, the so-called Reduced Binary Loopy Belief Propagation (RB-LBP), that significantly outperforms LBP in terms of scalability.

2. SCF PROBLEMS AS FACTOR GRAPHS

In LBP the SCF problem is casted into a factor graph composed of variables and factors. A single variable is created for each participant in the SC. The values (states) of each participant's variable encode the individual decisions that the agent needs to make regarding her exchange relationships plus an inactive state. For example, say that an agent needs to purchase a good to produce another one. Consider also that there are three possible producers for the requested good and three possible consumers for the produced good. Therefore, the agent's variable will have 10 states. That is, one for each of the producer-consumer combinations plus an inactive state. Notice that the number of states of an agent's variable grows exponentially with the number of agents and goods.

Agents' buying and selling prices are introduced by means of **activation factors**. Each agent has an activation factor that stores a value of zero whenever the agent is inactive and the agent's buying or selling price otherwise. Furthermore, in the factor graph, variables corresponding to potential partners are connected through a **compatibility factor**. Each of these factors encodes the compatibility between the decisions of the two agents involved. Two agents' decisions are incompatible whenever one of them is willing to trade with the other, but the other does not. Notice that the size of the compatibility factors is the product of the sizes of the variables it connects. Therefore, the memory needed by an agent to store factors grows exponentially to the number of agents and goods. Moreover, the messages exchanged between two agents encode their preferences over each other states. As a consequence, the communication requirements of LBP are also exponential to the number of goods and agents.

[1]We address the reader to [2] for a description of max-sum.

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

Measure	LBP	RB-LBP
Memory needed per agent to store the preferences over her state	$\mathcal{O}(A^G)$	$\mathcal{O}(G \cdot A)$
Size of largest factor	$\mathcal{O}(A^{2G})$	$\mathcal{O}(1)$
Maximum memory needed per agent (to store both preferences and factors)	$\boldsymbol{\mathcal{O}(G \cdot A^{2G+1})}$	$\boldsymbol{\mathcal{O}(G \cdot A)}$
Maximum message size	$\mathcal{O}(A^G)$	$\mathcal{O}(1)$
Maximum bandwith consumed per agent and iteration	$\boldsymbol{\mathcal{O}(G \cdot A^{G+1})}$	$\boldsymbol{\mathcal{O}(G \cdot A)}$
Overall consumed bandwith	$\mathcal{O}(n \cdot G \cdot A^{G+1})$	$\mathcal{O}(n \cdot G \cdot A)$

Table 1: Required resources: LBP vs. RB-LBP.

2.1 Scaling up supply chain formation

In order to cope with the scalability issues of LBP, we model the SCF problem as a binary factor graph containing only binary variables.

In this new model, each agent is aware of two sets of variables that encode her decisions to collaborate with potential partners. On the one hand, each agent encodes whether she is active (part of the SC) or not by means of an **activation variable**. On the other hand, each agent encodes her decision to trade a particular good with a particular producer/consumer using an **option variable**. Notice that the number of variables an agent needs to encode her decisions is linear to the number of possible exchanges she is involved in.

First, to guarantee that only one of the providers of a given good is selected, we make use of a **selection factor**. A selection factor links the activation variable from the agent with the option variables for that good. Second, we need to guarantee that the decisions from different agents are coherent among them. Thus, we add an **equality factor** constraining the seller's option variable and the buyer's option variable to be either both 1 or both 0. Notice that there is no need to store selection and activation factors in memory since they can be encoded as logical expressions.

Then, we show that, since we only employ binary variables and hard constraints, we can greatly reduce the computation required to assess messages. First, we only consider the configurations of variables that satisfy equality and selection factors. Second, instead of sending messages with a value for each of the two states of each variable, in RB-LBP messages contain the difference between these two values. Both changes together severely reduce the computation needed to assess messages. Moreover, since each agent only exchanges a single value with each of her neighbours, bandwidth requirements in RB-LBP scale linearly with the number of goods and agents.

Worst case memory and bandwidth requirements for both RB-LBP and LBP are summarized in table 1. A denotes the maximum number of agents connected to a good, G denotes the maximum number of goods an agent is interested in, and n stands for the total number of agents in the network.

3. EVALUATION

We benchmarked RB-LBP against LBP in the networks described by Walsh et al. in [4] and in larger networks with higher degrees of competition (in terms of number of providers offering each good). In the networks described by Walsh et al., RB-LBP requires from 2 up to 13 times less memory than LBP depending on the network structure. Moreover, the bandwidth consumed by an agent during an LBP iteration is up to 5 times larger than RB-LBP's.

For larger networks (up to 500 agents and 50 goods), LBP memory requirements are up to 5 orders of magnitude greater than for RB-LBP. Bandwidth usage for LBP is up to 787 times larger than for RB-LBP and, regarding computational time, RB-LBP is up to 20 times faster than LBP. Finally, the median SC value obtained by RB-LBP is up to 2 times greater than those obtained by LBP.

4. CONCLUSIONS AND FUTURE WORK

In this paper we have described RB-LBP, a novel approach for decentralized SCF. We have shown both theoretically and experimentally that RB-LBP scales nicely to market scenarios with larger number of participants and increasing competition. Our experimental results show that RB-LBP can significantly reduce the usage of memory and communication several orders of magnitude with respect to LBP. Furthermore, RB-LBP produces up to two times higher value supply chains and has smaller time complexity. Therefore, RB-LBP allows to tackle large-scale decentralized SCF problems.

Up to date approaches for decentralized SCF [4, 6] can only be applied to networks where agents can produce at most a single good. In order to compare with previously existing approaches, all the experimental results in this paper are over this kind of networks. However, RB-LBP can readily be applied to scenarios where producers can deliver more than one good. Experimentally evaluating RB-LBP in these scenarios and over a variety of actual-world network structures is left as future work.

5. ACKNOWLEDGMENTS

This work has been funded by projects EVE (TIN2009-14702-C02-01), AT (CSD2007-0022), CSIC 201050I008, and the Generalitat of Catalunya (2009-SGR-1434).

6. REFERENCES

[1] J. Cerquides, U. Endriss, A. Giovannucci, and J. A. Rodriguez-Aguilar. Bidding languages and winner determination for mixed multi-unit combinatorial auctions. In *IJCAI*, pages 1221–1226. Morgan Kaufmann Publishers Inc., 2007.

[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 (2)*, pages 639–646, 2008.

[3] V. Fionda and G. Greco. Charting the tractability frontier of mixed multi-unit combinatorial auctions. In C. Boutilier, editor, *IJCAI*, pages 134–139, 2009.

[4] W. E. Walsh and M. P. Wellman. Decentralized supply chain formation: A market protocol and competitive equilibrium analysis. *J. Artif. Intell. Res. (JAIR)*, 19:513–567, 2003.

[5] W. E. Walsh, M. P. Wellman, and F. Ygge. Combinatorial auctions for supply chain formation. In *2nd ACM Conference on Electronic Commerce*, EC '00, pages 260–269, New York, NY, USA, 2000. ACM.

[6] M. Winsper and M. Chli. Decentralised supply chain formation: A belief propagation-based approach. In *ECAI*, pages 1125–1126, Amsterdam, The Netherlands, The Netherlands, 2010. IOS Press.