In *Lecture Notes in Artificial Intelligence, vol 1788: Agent Mediated Electronic Commerce II,*
Springer-Verlag, 2000

Modeling Supply Chain Formation in Multiagent Systems*

William E. Walsh Michael P. Wellman

Artificial Intelligence Laboratory
University of Michigan
1101 Beal Avenue, Ann Arbor, MI 48109-2110 USA
{wew, wellman}@umich.edu

Abstract. Supply chain formation is an important problem in the commercial world, and can be improved by greater automated support. Hence, the multiagent systems community should work to develop new solutions to the problem. The problem is complex and challenging, and a complete model must encompass a number of issues. In this paper we highlight some issues that must be understood to make progress in modeling supply chain formation.

1 Introduction

Much of the popular coverage of electronic commerce has focused on technology for facilitating bilateral exchange between customers and merchants. However, complex economic activity often involves interrelated exchange relationships among multiple levels of production, often referred to as a ***supply chain***. Whereas a great deal of current commercial effort is being devoted to technology to support and manage supply chains, much of it is oriented toward maintaining pre-existing relationships in the chain. To achieve the oft-expressed visions of dynamically forming and dissolving business interactions (e.g., the rhetoric of "virtual corporations") requires automated support for ***supply chain formation***, the process of bottom-up assembly of complex production and exchange relationships.

Automated support can extend beyond speeding the communications, calculations, and routine computation in the supply chain formation process. The artificial intelligence and multiagent system (MAS) communities are well-positioned to develop technology that will increasingly automate the decision making in business interactions. However, we must recognize that merely describing computational entities as "agents" will not immediately solve the problem. Institutions and modes of interaction must be carefully designed to achieve desirable behavior. Supply chain formation is a complex and challenging problem, and a complete model would encompass a number of issues. We focus on the issue of resource contention, which has not been well addressed in most models relevant to supply chain formation.

In this paper we highlight some issues that must be understood to make progress in modeling supply chain formation. We outline these issues broadly in Section 2. In Section 3 we discuss several relevant models in multiagent systems, and highlight task dependency networks for representing resource contention. In Section 4 we describe some difficulties that arise from resource contention, and suggest that market-based approaches can be effective in solving them. In Section 5 we identify some open problems in supply chain formation, and conclude in Section 6.

2 Aspects of Supply Chain Formation

The defining characteristic of supply chain formation is ***hierarchical subtask decomposition***. Agents have specialized capabilities and can perform only certain combinations of tasks, or produce certain resources. In order to complete a complex task, an agent may delegate subtasks to other agents, which may in turn delegate further subtasks.

Resource contention constrains the set of feasible supply chains. When agents require a common resource (e.g., a subtask achievement, or something tangible like a piece of equipment) to complete their own tasks, resource scarcity may exclude these agents from jointly operating in the supply chain.

Agents aim to maximize their ***value*** in a supply chain. The model should account for values over multiple attributes such as monetary cost, time, and quality.

In a MAS containing autonomous, self-interested agents with private information, no entity has all information necessary to centrally form supply chains. This ***decentralization*** constraint implies that we should distribute decision making throughout the supply chain. Ideally, agents would be able to form supply chains in a strictly bottom-up manner, requiring only local communications and limited knowledge of the rest of the system.

Even when distributed, agent interactions during supply chain formation can be complex. Agents do not generally have the incentive to truthfully reveal information, hence we must analyze ***strategic interactions*** between agents.

Uncertainty introduces many complications in a model of supply chain formation. The constructed supply chain can be disrupted by failures in agents, failure in communications, and intentional fraud. Agents may be allowed to legally decommit from their contracts for a cost (e.g. to take advantage of better opportunities). The negotiation process itself may be affected by uncertain events. Bounds on rationality, variations in knowledge across agents, and multiplicity of equilibria make an agent's negotiation strategy uncertain. Failures in agents, negotiation mediators, and communications channels can occur during negotiation. Even if we know the full negotiation protocol and agent strategies and no failures occur, asynchronous communications introduces random elements into negotiation.

3 Models of Task Allocation and Supply Chains

Several existing models from the MAS research literature capture some, but not all, elements of supply chain formation. We review some of the dominant models of relevance here.

Rosenschein and Zlotkin [6] study a class of problems they call Task Oriented Domains (TODs). A TOD consists of a set of tasks that must be accomplished, and a (generally non-additive) cost function over sets of tasks. Sandholm [7] examines a generalization of TODs to include agent-dependent costs and constraints on task achievement. Agents are initially allocated a set of tasks, and may negotiate mutually beneficial exchanges of tasks to lower their costs. Andersson and Sandholm extend the model to allow agents to pay to decommit from contracts when better contracts become available [1]. These models address strategic interactions and increasing value via bilateral and multilateral task exchange, but do not incorporate the subtasking relationships characteristic of supply chains.

Hierarchical subtask decomposition is a core feature of the model underlying the CONTRACT NET protocol [2]. Subtasking proceeds in a distributed manner, in that agents communicate only with potential and actual subtasking partners. The protocol can flexibly accommodate multiattribute optimization criteria. However, these optimization criteria are applied locally within a level, in a greedy fashion, and do not reflect negotiation at other levels in the evolving supply chain.

Joshi et al. [3] consider issues arising from simultaneous negotiation of multiple subtasking issues at various levels of a supply chain. In their asynchronous model, agents may have the opportunity to finalize a contract while other negotiations are still pending. This uncertainty induces a complex decision problem for agents that do not wish to overextend their commitments.

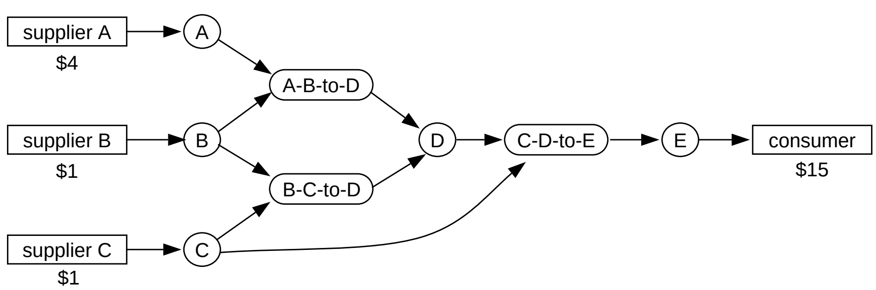

Fig. 1. A task dependency network.

Balancing interacting subtask commitments at multiple levels is especially important in the presence of resource contention. We represent resource contention in a supply chain with a ***task dependency network*** [8]. The nodes of a network represent agents and goods. We use the term ***good*** to refer to any task or discrete resource provided or needed by agents. The edges in the network represent input/output relationships between agents and goods. An edge from good to an agent indicates that the agent can make use of one unit of the good, and an edge from an agent to a good indicates that the agent can provide one unit of the good. A ***consumer*** is an agent that obtains some value from acquiring a good. A ***producer*** can provide one unit of one output good, contingent on acquiring each of a set of input goods. A ***supplier*** can provide a good, at some cost, without requiring any inputs.

Figure 1 shows an example task dependency network. The goods are indicated by circles, the consumer and suppliers by boxes, and the producers by curved boxes. The consumer value and supplier costs are indicated under the respective agents.

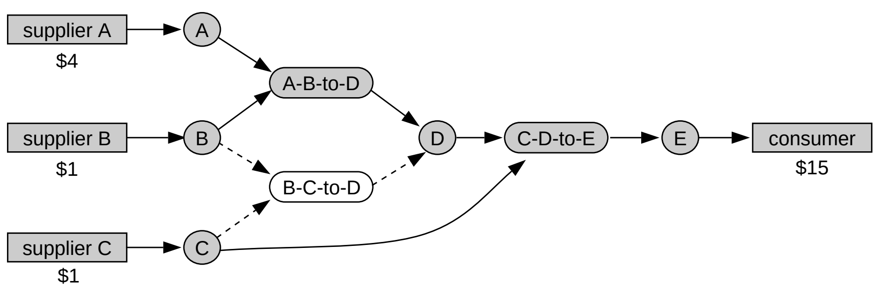

Fig. 2. The optimum allocation (and only solution) in the network.

An ***allocation*** is a subgraph containing all agents that acquire or provide goods, all goods that are acquired or provided, and all edges corresponding to the acquiring/providing relationships between the goods and agents in the allocation. A ***feasible allocation*** is one in which: 1) producers that provide their outputs also acquire all necessary inputs, and 2) all goods are in material balance (the number of edges into a good equals the number of edges out of a good). An ***optimal allocation*** is a feasible allocation that maximizes the difference between the sum of consumer value, and sum of supplier cost in the allocation. A ***solution*** is a feasible allocation in which one or more consumers acquire a desired good. Figure 2 shows the optimum allocation (and only solution) for the network from Figure 1. The allocation includes the shaded agents and goods and the solid edges.

4 Problems With Resource Contention

Consider a simple greedy allocation protocol, whereby subtask requests flow from consumers to suppliers, and subtask offers flow back to the consumer in a single pass. At each point in the offer flow, an agent sends offers as a function of the best offers it receives. Figures 1 and 2 show an example of how such a greedy policy could run into trouble with resource contention. Locally, it appears that B-C-to-D could provide D cheaper than A-B-to-D. However, the only solution includes C-D-to-E, which requires the one available unit of C. This necessarily excludes B-C-to-D from the solution.

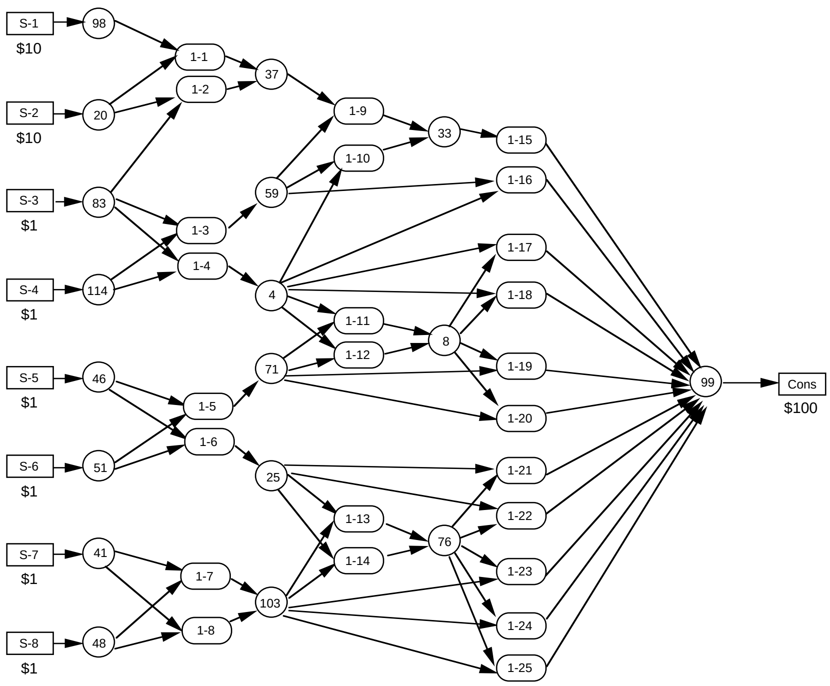

Fig. 3. A harder task dependency network.

We might address this problem by augmenting a greedy protocol with lookahead or backtracking capabilities to identify these constraints. However, we argue that even with these capabilities, the protocol would flounder when faced with more complicated networks such as the one in Figure 3. Any solution must include suppliers S-1 and S-2 (the optimum is shown in Figure 4). However, because these suppliers have the highest costs, the network would severely mislead a greedy protocol.

We can show that the problem of finding solutions is NP-complete, hence it should not be surprising that greedy approaches would have difficulty. Indeed, we should not expect any decentralized protocol to always form optimal supply chains.

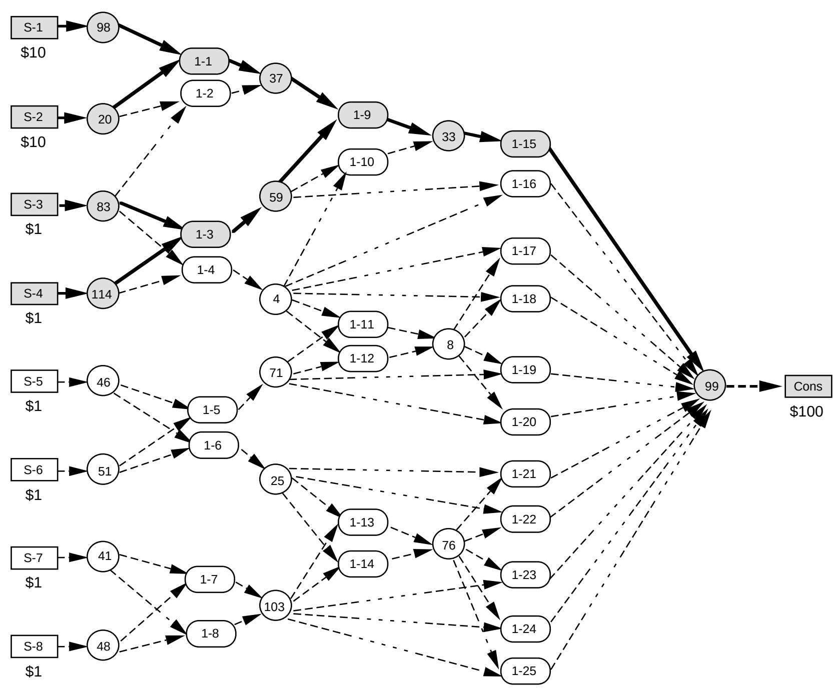

Fig. 4. The optimum allocation (and only solution) to the harder network.

We identify a market-based protocol that, while not uniformly optimal, performs well on task dependency networks [8]. Markets utilize ***price systems*** to guide agent decisions in a decentralized manner. Agents negotiate through simultaneous, ascending auctions, one for each good. We examine simple, nonstrategic agent bidding policies, requiring only local information. The protocol quickly finds solutions for the example problems shown in this paper. Experiments suggest that the protocol reliably converges to solutions (when they exist, and when the consumers have sufficiently high value) [10]. When combined with contract decommitment to remove dead ends in the supply chain, the protocol produces high value allocations on average [9].

5 Open Problems

Task dependency networks account for resource contention in supply chain formation, which has not been adequately addressed in other models. We have identified a market-based protocol that performs well on task dependency networks, but there remain many open problems to be addressed.

Strategic interactions are extremely complicated in any reasonable model of supply chain formation. In our work we have assumed simple agent bidding policies that work well from a system perspective. We must perform a strategic analysis to determine the plausibility of the policies, and how they fare against other policies. Additionally, we would like to understand how agents might exploit knowledge of the network to their advantage.

To a certain degree, we can model multiattribute value in task dependency networks. For instance, we can model discrete-time scheduling preferences by replicating agents over periods in a time interval. However, this replication becomes infeasible as the number of attributes grows. Incorporating multiattribute valuations more directly would require an extension to the market protocol.

In task dependency networks, we can also model flexible production of multiple units and multiple output goods by representing a single production entity as multiple producer agents, one for each production choice. Again, this can be reasonable for small-scale multiplication, but may not scale well with extensive production choices. Moreover, we cannot claim that the multiple agents are strategically equivalent to the single entity they represent. We consider these to be interesting challenges, important to future development of our model.

In a highly dynamic system we might expect changes to necessitate reallocation after quiescence. An adaptive system would detect—in a decentralized fashion—when lost resources cause infeasibility or when new opportunities make an improved allocation possible. We plan to study how well market allocations degrade or improve in response to reallocations, and how to compensate agents that lose value in the reallocation.

Rational agents would use probabilistic reasoning to adjust their negotiations to account for the possibility of later reallocation. We can also introduce devices such as securities and contingency contracts to improve allocations in the face of uncertainty.

Although we have presented the task dependency networks as pre-defined entities, in a dynamic system we would expect the goods of interest and the available auctions to change over time. Mullen and Wellman [4] studied mechanisms and policies for starting and maintaining auctions, and others [5, 11] are working to design *goods description languages* to allow agents to communicate their needs.

6 Conclusion

Supply chain formation is an important problem in the commercial world, and can be improved by greater automated support. The problem is salient to the MAS community and deserving of continued research.

Acknowledgments This work was supported by a NASA/Jet Propulsion Laboratory Graduate Student Researcher fellowship and DARPA grant F30602-97-1-0228 from the Information Survivability program.

References

[1] Martin R. Andersson and Tuomas W. Sandholm. Leveled commitment contracting among myopic individually rational agents. In *Third International Conference on Multi-Agent Systems*, pages 26–33, 1998.

[2] Randall Davis and Reid G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20:63–109, 1983.

[3] Pawan Joshi, Mayuresh Oke, Vivek Sharma, and Dharmaraj Veeramani. Issues in dynamic and highly-distributed configuration of supply webs. In *First IAC Workshop on Internet-Based Negotiation Technologies*, 1999.

[4] Tracy Mullen and Michael P. Wellman. The auction manager: Market middleware for large-scale electronic commerce. In *Third USENIX Workshop on Electronic Commerce*, pages 37–47, Boston, 1998.

[5] Daniel M. Reeves, Benjamin N. Grosof, Michael P. Wellman, and Hoi Y. Chan. Automated negotiations from formal contract descriptions (work-in-progress report). In *First IAC Workshop on Internet-Based negotiation Technologies*, March 1999.

[6] Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter*. MIT Press, 1994.

[7] Tuomas W. Sandholm. An implementation of the CONTRACT NET protocol based on marginal cost calculations. In *Eleventh National Conference on Artificial Intelligence*, pages 256–262, 1993.

[8] William E. Walsh and Michael P. Wellman. A market protocol for decentralized task allocation. In *Third International Conference on Multi-Agent Systems*, pages 325–332, 1998.

[9] William E. Walsh and Michael P. Wellman. Efficiency and equilibrium in task allocation economies with hierarchical dependencies. In *Sixteenth International Joint Conference on Artificial Intelligence*, 1999.

[10] William E. Walsh, Michael P. Wellman, Peter R. Wurman, and Jeffrey K. MacKie-Mason. Some economics of market-based distributed scheduling. In *Eighteenth International Conference on Distributed Computing Systems*, pages 612–621, 1998.

[11] Peter Weinstein and William P. Birmingham. Service classification in a proto-organic society of agents. In *IJCAI-97 Workshop on Artificial Intelligence in Digital Libraries*, 1997.