Protocol/Mechanism Design for Cooperation/Competition

Makoto Yokoo
Department of Intelligent Systems
Graduate School of Information Science and Electrical Engineering
Kyushu University
6-10-1 Hakozaki, Higashi-ku, Fukuoka, 812-8581 Japan
e-mail: yokoo@is.kyushu-u.ac.jp
url: lang.is.kyushu-u.ac.jp/˜yokoo/

Abstract

Developing interaction rules/protocols among multiple agents is one of the central research topics in multi-agent systems. For cooperative agents, we need to develop protocols so that agents can achieve some common goal if they follow the protocol. Also, for competitive/selfish agents, we need to design mechanisms/protocols so that some socially desirable outcome can be achieved, even if agents act selfishly. This article presents a brief overview of the author's works on this topic over the last five years.

1. Protocol Design for Cooperative Agents

1.1. Distributed Constraint Satisfaction

When there are multiple agents in a shared environment, there usually exist constraints among the possible actions of these agents. A distributed constraint satisfaction problem (distributed CSP) is a problem to find a consistent combination of actions that satisfies these inter-agent constraints [34, 35]. The research on constraint satisfaction problems has a long and distinguished history in AI as a general framework that can formalize various application problems [14]. Similarly, a distributed CSP is a fundamental problem for achieving coordination among agents and can formalize various application problems in multi-agent systems.

A typical example of a CSP is a puzzle called 8-queens (Figure 1). The objective is to place eight queens on a chess board (8×8 squares) so that these queens will not threaten each other. This problem is called a constraint satisfaction problem since the objective is to find a configuration that satisfies the given conditions (constraints).

Formally, a CSP consists of n variables $x_1, x_2, \ldots, x_n$, whose values are taken from finite, discrete domains $D_1, D_2, \ldots,\ D_n$, respectively, and a set of constraints on their values. Solving a CSP is equivalent to finding the assignment of values to all variables such that all constraints are satisfied. Since constraint satisfaction is NP-complete in general, a trial-and-error exploration of alternatives is inevitable.

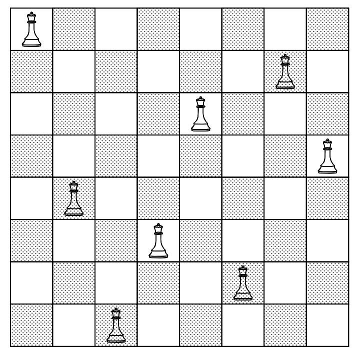

Figure 1. Example of a constraint satisfaction problem (8-queens)

A distributed CSP is a CSP in which the variables and constraints are distributed among automated agents. Each agent has some variables and tries to determine their values. However, there also exist inter-agent constraints, and the value assignment must satisfy these inter-agent constraints.

Various application problems in multi-agent systems that involve finding a consistent combination of agent actions can be formalized as distributed CSPs. One example is a distributed resource allocation task, such as a distributed sensor network [13], a distributed resource allocation problem in a communication network described in [2], or a channel assignment problem in a cellular radio network [38].

In these problems, each agent has its own tasks, and there are several ways (plans) to perform each task. Since resources are shared among agents, there exist con-

AAMAS'04, July 19-23, 2004, New York, New York, USA.

straints/contention between plans. The goal is to find the combination of plans that enables all the tasks to be executed simultaneously. Many other application problems that involve finding a consistent combination of agent actions/decisions (e.g., distributed scheduling and distributed interpretation problems) can be formalized as distributed CSPs.

In our earlier works, we developed a series of algorithms for solving distributed CSPs, including a basic backtracking algorithm called asynchronous backtracking, a more efficient algorithm called asynchronous weak-commitment search algorithm, in which the order of agents/variables is changed dynamically, and an iterative improvement algorithm called the distributed breakout algorithm [32, 34, 35, 36, 37].

The last five years have seen significant advances in various aspects of distributed CSPs. For example, we have shown that by introducing nogood learning, i.e., resolving/recording effective nogoods, the performance of the asynchronous weak-commitment search algorithm (as well as the asynchronous backtracking algorithm) can be significantly improved so that they can be comparable to the distributed breakout algorithm [6]. Furthermore, we have proposed an algorithm that is drastically different from traditional algorithms, i.e., it utilizes a market mechanism for solving satisfiability problems (SAT) [31, 30]. Also, we have examined a method for handling the case where the local problem of each agent is complex [7].

In many real-life applications, problems are over-constrained, i.e., there exists no solution that satisfies all constraints completely. In that case, we need to settle for an incomplete solution. We have proposed the formalization of hierarchical distributed CSPs, in which constraints are ordered by their relative importance [5]. If a problem is over-constrained, we are going to give up less important constraints. Furthermore, we have developed an algorithm called Asynchronous Distributed OPtimization (ADOPT) that can solve more general distributed constraint optimization problems [18, 19]. In this algorithm, each agent acts concurrently and asynchronously as in the asynchronous backtracking algorithm. Each agent keeps on increasing the threshold for backtracking. Then, they eventually find an optimal solution.

A major motivation for solving a distributed CSP without gathering all information into one server is the concern about privacy/security. However, existing distributed CSP algorithms leak some information during the search process, and privacy/security issues are not dealt with formally. To alleviate this problem, we have developed an algorithm that utilizes a public key encryption scheme so that the algorithm does not leak information, i.e., agents cannot obtain any additional information on the value assignment of variables that belong to other agents [49].

Many researchers have now started working on distributed CSP. There have been workshops on distributed CSP at CP-2000 (Int. Conf. on Principles and Practice of Constraint Programming), IJCAI-2001, AAMAS-2002, and IJCAI-2003. The next one will be at CP-2004.

1.2. Multi-agent POMDP

Partially Observable Markov Decision Processes (POMDPs) are emerging as a popular approach for modeling multi-agent teamwork where a group of agents work together to jointly maximize a reward function [22] This framework is quite important since it is theoretically well founded and can be used for re-formalizing existing models of multi-agent systems.

Even if the underlying system can be modeled as a POMDP, if there exist multiple agents, then the system including other agents is no longer POMDP from a single agent perspective. We have developed an innovative way to represent the belief state of an agent. By using this representation, we can model the system including other agents as a POMDP, given that the policies of other agents are fixed [21].

By utilizing this characteristic, we can efficiently find an optimal policy (i.e., a best response) for an agent by applying dynamic programming [21]. Also, this representation can precisely define the effect of communication; thus we can also find an optimal policy when communication requires certain costs [20].

2. Mechanism Design for Competitive Agents

2.1. False-name Bids

Internet auctions have become an especially popular part of Electronic Commerce (EC). Various theoretical and practical studies on Internet auctions have already been conducted. Among these studies, those on combinatorial auctions have lately attracted considerable attention ([3] is a good survey article). Although conventional auctions sell a single good at a time, combinatorial auctions sell multiple goods with interdependent values simultaneously and allow the bidders to bid on any combination of goods. In a combinatorial auction, a bidder can express complementary/substitutable preferences over multiple goods. By taking into account such preferences, economic efficiency can be enhanced.

Although the Internet provides an excellent infrastructure for executing combinatorial auctions, we must consider the possibility of new types of cheating. For example, a bidder may try to profit from submitting false bids under fictitious names such as multiple e-mail addresses. Such an

AAMAS'04, July 19-23, 2004, New York, New York, USA.

action is very difficult to detect since identifying each participant on the Internet is virtually impossible. We call a bid made under a fictitious name a *false-name bid*. Also, we say a protocol is *false-name-proof* if truth-telling without using false-name bids is a dominant strategy for each bidder. This is a natural extension of the traditional definition of strategy-proofness.

The problems resulting from collusion have been discussed by many researchers [15, 16, 17]. Compared with collusion, a false-name bid is easier to execute on the Internet since obtaining additional identifiers, such as another e-mail address, is cheap. We can consider false-name bids as a very restricted subclass of collusion.

In [39, 45], we have analyzed the effects of false-name bids on combinatorial auction protocols. The obtained results can be summarized as follows.

- The Vickrey-Clarke-Groves (VCG) mechanism [1, 4, 29], which is strategy-proof and Pareto efficient if there exists no false-name bid, is not false-name-proof.
- There exists no false-name-proof combinatorial auction protocol that satisfies Pareto efficiency.
- We identify one sufficient condition where the VCG mechanism is false-name-proof, i.e., a surplus function is *concave* over bidders.

Also, we have developed a series of protocols that are false-name-proof in various settings: a combinatorial auction protocol called the Leveled Division Set (LDS) protocol [40, 41, 42], multi-unit auction protocols [11, 12, 28, 44], and double auction protocols [23, 24, 43, 46].

Furthermore, we have identified a distinctive class of combinatorial auction protocols called a Price-oriented, Rationing-free (PORF) protocol, which can be used as a guideline for developing strategy/false-name proof protocols [33]. The outline of a PORF protocol is as follows: (i) for each bidder, the price of each bundle of goods is determined independently of his/her own declaration (while it can depend on the declarations of other bidders), (ii) we allocate each bidder a bundle that maximizes his/her utility independently of the allocations of other bidders (i.e., rationing-free).

Although a PORF protocol appears quite different from traditional protocol descriptions, surprisingly, it is a sufficient and necessary condition for a protocol to be strategy-proof. Furthermore, we show that a PORF protocol satisfying additional conditions is false-name-proof; at the same time, any false-name-proof protocol can be described as a PORF protocol that satisfies the additional conditions. A PORF protocol is an innovative characterization of strategy-proof protocols and the first attempt to characterize false-name-proof protocols. Such a characterization is not only theoretically significant but also useful in practice, since it can serve as a guideline for developing new strategy/false-name proof protocols. We have developed a new false-name-proof protocol based on the concept of a PORF protocol.

2.2. Secure Protocols

We developed secure dynamic programming protocols that utilize information security techniques [26, 47]. By using this method, multiple agents can solve a combinatorial optimization problem among them (e.g., winner-determination in combinatorial auctions) without leaking their private information to other agents. More specifically, in these methods, multiple servers cooperatively perform dynamic programming procedures for solving a combinatorial optimization problem by using the private information sent from agents as inputs. Although the servers can compute the optimal solution correctly, the inputs are kept secret from these servers. Such a secure protocol is important when a fully trusted agent is not available, e.g., an auctioneer cannot be fully trusted in a combinatorial auction.

Furthermore, we have developed a protocol that can perform the Vickrey-Clarke-Groves (VCG) mechanism for a very general setting [27]. Also, even if a protocol itself is strategy-proof, if the protocol is executed by utilizing a P2P network, which does not have a centralized server, then participants might have an incentive to deviate from the protocol [25]. We have developed a secure VCG-type protocol that does not require third-party servers, i.e., agents can execute the protocol by themselves [48].

2.3. Auctions under Asymmetric Information

In Internet auctions, it is often difficult to determine the quality of auctioned goods. We consider a situation where some experts can judge the quality of auctioned goods correctly, while many armatures cannot do so. We have developed a series of protocols that gives experts an incentive to truthfully declare the quality; thus we can achieve Pareto efficient allocations in many cases [8, 9, 10].

Acknowledgements

The results presented in this article are joint works with many people and I'd like to thank my co-authors. In particular, I'd like to express my sincere thanks to Toru Ishida and Edmund H. Durfee, who have been my mentors. Also, I'm grateful to my longtime co-researcher Katsutoshi Hirayama. Furthermore, I'd like to thank my former colleagues at NTT Laboratories: Yuko Sakurai, Shigeo Matsubara, Koutaro Suzuki, Atsushi Iwasaki, Kenji Terada, and Takayuki Suyama. Also, I visited various places and was inspired by working with many people. I'd like to thank them

AAMAS'04, July 19-23, 2004, New York, New York, USA.

all, in particular, Tuomas Sandholm, Boi Faltings, Marius C. Silaghi, Takayuki Ito, William Walsh, Michael P. Wellman, Pragnesh Jay Modi, Ranjit Nair, David Pynadath, and Milind Tambe.

References

[1] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 2:19–33, 1971.

[2] S. E. Conry, K. Kuwabara, V. R. Lesser, and R. A. Meyer. Multistage negotiation for distributed constraint satisfaction. *IEEE Transactions on Systems, Man and Cybernetics*, 21(6):1462–1477, 1991.

[3] S. de Vries and R. V. Vohra. Combinatorial auctions: A survey. *INFORMS Journal on Computing*, 15, 2003.

[4] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[5] K. Hirayama and M. Yokoo. An approach to over-constrained distributed constraint satisfaction problems: Distributed hierarchical constraint satisfaction. In *Proceedings of the Fifth International Conference on Multi-agent Systems (ICMAS-2000)*, 2000.

[6] K. Hirayama and M. Yokoo. The effect of nogood learning in distributed constraint satisfaction. In *Proceedings of the Twentieth International Conference on Distributed Computing Systems (ICDCS-2000)*, pages 169–177, 2000.

[7] K. Hirayama and M. Yokoo. Local search for distributed sat with complex local problems. In *Proceedings of the First International joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2002)*, pages 1199–1206, 2002.

[8] T. Ito, M. Yokoo, and S. Matsubara. Designing an auciton protocol under asymmetric information on nature's selection. In *Proceedings of the First International joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2002)*, pages 61–68, 2002.

[9] T. Ito, M. Yokoo, and S. Matsubara. Toward a combinatorial auction protocol among experts and amateurs:the case of single-skilled experts. In *Proceedings of the Second International joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2003)*, pages 481–488, 2003.

[10] T. Ito, M. Yokoo, and S. Matsubara. A combinatorial auction among versatile experts and amateurs. In *Proceedings of the Third International joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2004)*, 2004.

[11] A. Iwasaki, M. Yokoo, and K. Terada. A robust open ascending-price multi-unit auction protocol against false-name bids. In *Proceedings of the Fourth ACM Conference on Electronic Commerce (EC-2003)*, pages 85–92, 2003.

[12] A. Iwasaki, M. Yokoo, and K. Terada. A robust open ascending-price multi-unit auction protocol against false-name bids. *Decision Support Systems*, (to appear).

[13] V. Lesser, C. L. Ortiz, and M. Tambe, editors. *Distributed Sensor Networks: A Multiagent Perspective*. Kluwer, 2003.

[14] A. K. Mackworth. Constraint satisfaction. In S. C. Shapiro, editor, *Encyclopedia of Artificial Intelligence*, pages 285–293. John Wiley & Sons, New York, 1992.

[15] R. P. McAfee and J. McMillan. Auction and bidding. *Journal of Economic Literature*, 25:699–738, 1987.

[16] R. P. McAfee and J. McMillan. Bidding rings. *American Economic Review*, 82(3):579–599, 1992.

[17] P. R. Milgrom. Putting auction theory to work: The simultaneous ascending auction. *The Journal of Political Economy*, 108(2):245–272, 2000.

[18] P. J. Modi, W.-M. Shen, M. Tambe, and M. Yokoo. An asynchronous complete method for distributed constraint optimization. In *Proceedings of the Second International joint Conference on Autonomous Agents (AAMAS-2003)*, pages 161–168, 2003.

[19] P. J. Modi, W.-M. Shen, M. Tambe, and M. Yokoo. An asynchronous complete method for distributed constraint optimization. *Artificial Intelligence Journal*, 2004. (to appear).

[20] R. Nair, S. Marsella, M. Tambe, D. Pynadath, and M. Yokoo. Taming decentralized pomdps: Towards efficient policy computation for multiagent settings. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-2003)*, pages 705–711, 2003.

[21] R. Nair, M. Roth, M. Yokoo, and M. Tambe. Communication for improving policy computation in distributed pomdps. In *Proceedings of the Third International joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2004)*, 2004.

[22] D. Pynadath and M. Tambe. The communicative multiagent team decision problems:analyzing teamwork theories and models. *Journal of Artificial Intelligence Research*, 16:389–423, 2002.

[23] Y. Sakurai and M. Yokoo. An average-case budget-non-negative double auction protocol. In *First International joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2002)*, pages 104–111, 2002.

[24] Y. Sakurai and M. Yokoo. A false-name-proof double auction protocol for arbitrary evaluation values. In *Second International joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2003)*, pages 329–336, 2003.

[25] J. Shneidman and D. C. Parkes. Rationality and self-interest in peer to peer networks. In *Proceedings of Second International Workshop on Peer-to-Peer Systems (IPTPS-2003)*, 2003.

[26] K. Suzuki and M. Yokoo. Secure combinatorial auctions by dynamic programming with polynomial secret sharing. In *Proceedings of Sixth International Financial Cryptography Conference (FC-02)*, Lecture Notes in Computer Science 2357, pages 44–56. Springer, 2002.

[27] K. Suzuki and M. Yokoo. Secure Generalized Vickrey Auction using homomorphic encryption. In *Proceedings of Seventh International Financial Cryptography Conference (FC-03)*, 2003.

[28] K. Terada and M. Yokoo. False-name-proof multi-unit auction protocol utilizing greedy allocation based on approximate evaluation values. In *Proceedings of the Second International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2003)*, pages 337–344, 2003.

[29] W. Vickrey. Counter speculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[30] W. E. Walsh, M. Yokoo, K. Hirayama, and M. P. Wellman. On market-inspired approaches to propositional satisfiability. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-2001)*, pages 1152–1158, 2001.

[31] W. E. Walsh, M. Yokoo, K. Hirayama, and M. P. Wellman. On market-inspired approaches to propositional satisfiability. *Artificial Intelligence*, 144(1–2):125–156, 2003.

[32] M. Yokoo. *Distributed Constraint Satisfaction: Foundation of Cooperation in Multi-agent Systems*. Springer, 2001.

[33] M. Yokoo. The characterization of strategy/false-name proof combinatorial auction protocols: Price-oriented, rationing-free protocol. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, pages 733–739, 2003.

[34] M. Yokoo, E. H. Durfee, T. Ishida, and K. Kuwabara. Distributed constraint satisfaction for formalizing distributed problem solving. In *Proceedings of the Twelfth IEEE International Conference on Distributed Computing Systems*, pages 614–621, 1992.

[35] M. Yokoo, E. H. Durfee, T. Ishida, and K. Kuwabara. The distributed constraint satisfaction problem: formalization and algorithms. *IEEE Transactions on Knowledge and Data Engineering*, 10(5):673–685, 1998.

[36] M. Yokoo and K. Hirayama. Distributed breakout algorithm for solving distributed constraint satisfaction problems. In *Proceedings of the Second International Conference on Multi-Agent Systems*, pages 401–408. MIT Press, 1996.

[37] M. Yokoo and K. Hirayama. Algorithms for distributed constraint satisfaction: A review. *Autonomous Agents and Multi-agent Systems*, 3(2):189–211, 2000.

[38] M. Yokoo and K. Hirayama. Frequency assignment for cellular mobile systems using constraint satisfaction techniques. In *Proceedings of the IEEE Annual Vehicular Technology Conference (VTC-2000 Spring)*, 2000.

[39] M. Yokoo, Y. Sakurai, and S. Matsubara. The effect of false-name declarations in mechanism design: Towards collective decision making on the Internet. In *Proceedings of the Twentieth International Conference on Distributed Computing Systems (ICDCS-2000)*, pages 146–153, 2000.

[40] M. Yokoo, Y. Sakurai, and S. Matsubara. Robust combinatorial auction protocol against false-name bids. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-2000)*, pages 110–115, 2000.

[41] M. Yokoo, Y. Sakurai, and S. Matsubara. Bundle design in robust combinatorial auction protocol against false-name bids. In *Proceedings of 17th International Joint Conference on Artificial Intelligence (IJCAI-2001)*, pages 1095–1101, 2001.

[42] M. Yokoo, Y. Sakurai, and S. Matsubara. Robust combinatorial auction protocol against false-name bids. *Artificial Intelligence*, 130(2):167–181, 2001.

[43] M. Yokoo, Y. Sakurai, and S. Matsubara. Robust double auction protocol against false-name bids. In *Proceedings of the 21st International Conference on Distributed Computing Systems (ICDCS-2001)*, pages 137–145, 2001.

[44] M. Yokoo, Y. Sakurai, and S. Matsubara. Robust multi-unit auction protocol against false-name bids. In *Proceedings of 17th International Joint Conference on Artificial Intelligence (IJCAI-2001)*, pages 1089–1094, 2001.

[45] M. Yokoo, Y. Sakurai, and S. Matsubara. The effect of false-name bids in combinatorial auctions: New fraud in Internet auctions. *Games and Economic Behavior*, 46(1):174–188, 2004.

[46] M. Yokoo, Y. Sakurai, and S. Matsubara. Robust double auction protocol against false-name bids. *Decision Support Systems*, (to appear).

[47] M. Yokoo and K. Suzuki. Secure multi-agent dynamic programming based on homomorphic encryption and its application to combinatorial auctions. In *Proceedings of the First International Conference on Autonomous Agents and Multi-agent Systems (AAMAS-2002)*, pages 112–119, 2002.

[48] M. Yokoo and K. Suzuki. Secure Generalized Vickrey Auction without third-party servers. In *Proceedings of Eighth International Financial Cryptography Conference (FC-04)*, 2004.

[49] M. Yokoo, K. Suzuki, and K. Hirayama. Secure distributed constraint satisfaction: Reaching agreement without revealing private information. In *Proceedings of the Eighth International Conference on Principles and Practice of Constraint Programming (CP-2002)*, pages 387–401, 2002.