RESEARCH STATEMENT

DANIEL M. REEVES

A human faced with a new game digests the rules, considers the motives and capabilities of the other players, and formulates a strategy for playing. My research aims to automate that process. A *Strategy Generation Engine* is a system that advises on how to play a game given a formal description of the possible interactions of the participants. Automated strategy generation is important for multiple aspects of ecommerce:

(1) Assisting in the design of strategies for known mechanisms.
(2) Creating agents that can participate in new mechanisms or dynamically generated mechanisms.
(3) Evaluating mechanisms automatically.

The last is tantamount to computational mechanism design—an exciting area of future work in which I intend to apply the methodology I have developed in my thesis.

In addition to explicit negotiation and trading in electronic markets, there is also a strategic dimension to commerce activities such as matchmaking, resource finding, advertising, recommendation, contracting, and executing transactions [18, 25]. Likewise, other digital realms involve pivotal strategic relationships. Examples include peer-to-peer resource sharing [16], formation of coalitions, teams, or affinity groups [1, 22, 23], scientific sharing of large-scale instrumentation and other infrastructure [5, 6], coordination of activity within organizations [10, 17], mobile computing [8], and grid computing [7].

Interactions among self-interested agents are properly modeled as *games*. When games operate in the digital realm, there is often increased opportunity for *designing* interaction mechanisms. Thus, the economic theory of mechanism design has found increasing practical application in such far-flung policy areas as deregulation of energy markets [12] and operation of the nationwide medical residency match [21]. Recently, computer scientists have also recognized the importance of incentives in distributed systems, many enthusiastically adopting the framework of game theory [15] and mechanism design [4, 13].

1. Previous Work

My thesis work concerns the generation and selection of strategies, in particular for trading agents paricipating in various market mechanisms. Games that trading agents face typically involve private information, such as an agent's valuation for a good being bought or sold, and fine-grained action spaces, such as bids of arbitrary amounts of money. Such games are well modeled as infinite games of incomplete information. Generating trading agent strategies means creating a system that can read the description of such a mechanism and output strategies for participating agents.

To make this more concrete, consider an extremely simple auction mechanism: a First-Price Sealed-Bid Auction (FPSB). This is a game in which each agent has one piece of private information: its valuation for an indivisible good being auctioned (in the simplest case we take the common-knowledge type distribution to be i.i.d. $U[0,1]$). Each agent also has a continuum of possible actions: its bid amount. The payoff to an agent is its valuation minus its bid if its bid is highest, and zero otherwise (with ties broken by fair coin toss). In my thesis I present an algorithm that can solve this game. That is, it takes the game description—for the two-player case—and outputs the unique symmetric Bayes-Nash equilibrium. (In this case, for an agent with valuation v, the equilibrium strategy is to bid $v/2$.) Of course, the Nash equilibrium strategy for the particular case of FPSB was identified by auction theorists before computational game solvers existed [24, 11]. My algorithm applies to a class of games that includes the above example [19].

The above method is tractable only for quite simple games. For example, mechanisms that involve iterated bidding and multiple auctions are not likely to be amenable to analytic approaches. For such games, I present an *empirical game methodology* comprising the following broad phases:

(1) Generate a small set of candidate strategies. For many domains this must be done semi-manually.
(2) Construct via simulation a (partial, approximate) empirical payoff matrix for the simplification of the game restricted to the candidate strategies.
(3) Analyze (ideally, solve) the empirical game.
(4) Assess the quality of the solutions with respect to the underlying full game.

My thesis discusses approaches for step 1, presents techniques for speeding up step 2, compares existing techniques for step 3 in the context of symmetry, and gives procedures for achieving step 4 [27, 20, 9, 14, 3, 28]. We use small games with known equilibria such as FPSB to test these methods.

Later chapters of my thesis (4, 5, and 6) apply this methodology to two much larger and more realistic market games: Simultaneous Ascending Auctions (SAA), and the Trading Agent Competition (TAC) travel-shopping game [31]. TAC was created by our research group and first held in 2000. It was held for the sixth time in August 2005 in Edinburgh. The domain involves travel agents shopping for travel packages for a group of hypothetical clients with varying preferences over length of trip, hotel quality, and entertainment options. The shopping involves participating in dozens of simultaneous auctions of various kinds. For example, hotels are sold in multi-unit English ascending auctions while entertainment tickets are bought and sold in continuous double auctions (like the stock market). An agent's payoff is the total utility it achieves for its clients, minus its net expenditure.

SAA is a far simpler model that still captures a core strategic issue in TAC: bidding for complementary goods in concurrent auctions. To apply step 1 to SAA and TAC, I present classes of strategies based on market price prediction. In particular we consider self-confirming price predictions and Walrasian equilibrium prices [14]. Given a set of candidate strategies in SAA or TAC, we apply the subsequent steps of our empirical game methodology to recommend effective strategies in these domains [30, 2, 26, 29].

In conclusion, my thesis makes the following key contributions:

(1) An algorithm to compute best-response strategies in a class of 2-player, one-shot, infinite games of incomplete information.

(2) An empirical game methodology for applying game-theoretic analysis to much larger games than previously possible.
(3) Theoretical and experimental evidence of the efficacy of our methodology.
(4) A price-prediction approach to strategy generation in simultaneous auctions for complementary goods.
(5) Application of the above methods to find good strategies in complex games.

2. Future Directions

I am excited about applying my work on trading agent strategy generation to new domains. Extending my previous work on price prediction strategies for simultaneous ascending auctions with complementarities, I and colleagues at Michigan have begun to explore new classes of strategies in simultaneous auctions for goods that are substitutes. In the near term, I would like to continue this research thread, but, as demonstrated by our application of our empirical game methodology to the Trading Agent Competition, my strategy generation approach is applicable to wide varieties of very complex games. In particular, I would like to apply my strategy generation methods to the problem of bidding in combinatorial and multiattribute auctions.

Computational mechanism design is the problem of designing mechanisms (i.e., game rules and protocols) for computational agents under constraints such as fairness and efficiency when agents play Nash equilibrium strategies. When the designer is also constrained by factors such as agents entering and leaving dynamically or pre-existing protocol constraints such as iterative bidding then the mechanism design problem can be very difficult. While my thesis includes small examples of designed mechanisms, this is a wide open area to apply my work on strategy generators. Having developed the state of the art in finding strategies for complex games (and with additional automation of the strategy generation process in order to apply it programmatically to variations on market mechanisms) I will be in a unique position to advance the field of computational mechanism design. By improving fairness and efficiency of a wide variety of electronically mediated interaction, I strongly believe that such research has and will continue to yield substantial and tangible benefits for society.

References

[1] Christopher H. Brooks and Edmund H. Durfee. Congregation formation in multiagent systems. *Autonomous Agents and Multiagent Systems*, 7:145–170, 2003.

[2] Shih-Fen Cheng, Evan Leung, Kevin M. Lochner, Kevin O'Malley, Daniel M. Reeves, L. Julian Schvartzman, and Michael P. Wellman. Walverine: A Walrasian trading agent. *Decision Support Systems*, 39:169–184, 2005.

[3] Shih-Fen Cheng, Daniel M. Reeves, Yevgeniy Vorobeychik, and Michael P. Wellman. Notes on equilibria in symmetric games. In *AAMAS-04 Workshop on Game Theory and Decision Theory*, New York, 2004.

[4] Rajdeep K. Dash, Nicholas R. Jennings, and David C. Parkes. Computational mechanism design: A call to arms. *IEEE Intelligent Systems*, 18:40–47, 2003.

[5] Thomas A. Finholt. Collaboratories as a new form of scientific organization. *Economics of Innovation an New Technology*, 12:5–25, 2003.

[6] Thomas A. Finholt and Gary M. Olson. From laboratories to collaboratories: A new organizational form for scientific collaboration. *Psychological Science*, 8:28–36, 1997.

[7] Ian Foster, Carl Kesselman, and Steven Tuecke. The anatomy of the grid: Enabling scalable virtual organizations. *International Journal of High Performance Computing Applications*, 15(3):200–222, August 2001.

[8] Eric J. Friedman and David C. Parkes. Pricing WiFi at Starbucks—issues in online mechanism design. In *Fourth ACM Conference on Electronic Commerce*, pages 240–241, San Diego, 2003.

[9] Jeffrey K. MacKie-Mason, Anna Osepayshvili, Daniel M. Reeves, and Michael P. Wellman. Price prediction strategies for market-based scheduling. In *Fourteenth International Conference on Automated Planning and Scheduling*, pages 244–252, Whistler, BC, 2004.

[10] Thomas W. Malone. Modeling coordination in organizations and markets. *Management Science*, 33:1317–1332, 1987.

[11] R. Preston McAfee and John McMillan. Auctions and bidding. *Journal of Economic Literature*, 25:699–738, 1987.

[12] Paul R. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.

[13] Noam Nisan and Amir Ronen. Algorithmic mechanism design. In *Thirty-first ACM Symposium on Theory of Computing*, pages 129–140, 1999.

[14] Anna Osepayshvili, Michael P. Wellman, Daniel M. Reeves, and Jeffrey K. MacKie-Mason. Self-confirming price prediction for bidding in simultaneous ascending auctions. In *Twenty-first Conference on Uncertainty in Artificial Intelligence*, pages 441–449, July 2005.

[15] Christos Papadimitriou. Algorithms, games, and the internet. In *Thirty-third Annual ACM Symposium on the Theory of Computing*, Hersonissos, Greece, 2001.

[16] Ilya Mironov Mark Lillibridge Philippe Golle, Kevin Leyton-Brown. Incentives for sharing in peer-to-peer networks. In *Second International Workshop on Electronic Commerce (WELCOM)*, pages 75–87, 2001.

[17] David V. Pynadeth and Milind Tambe. The communicative multiagent team decision problem: Analyzing teamwork theories and models. *Journal of Artificial Intelligence Research*, 16:389–423, 2002.

[18] Daniel M. Reeves, Benjamin N. Grosof, Michael P. Wellman, and Hoi Y. Chan. Automated negotiation from declarative contract descriptions. *Computational Intelligence*, 18:482–500, 2002.

[19] Daniel M. Reeves and Michael P. Wellman. Computing best response strategies in infinite games of incomplete information. In *Twentieth Conference on Uncertainty in Artificial Intelligence*, pages 470–478, Banff, 2004.

[20] Daniel M. Reeves, Michael P. Wellman, Jeffrey K. MacKie-Mason, and Anna Osepayshvili. Exploring bidding strategies for market-based scheduling. *Decision Support Systems*, 39(1):67–85, March 2005.

[21] Alvin E. Roth and Elliott Peranson. The redesign of the matching market for American physicians: Some engnineering aspects of economic design. *American Economic Review*, 89:748–780, 1999.

[22] Tuomas W. Sandholm and Victor R. Lesser. Coalitions among computationally bounded agents. *Artificial Intelligence*, 94:99–137, 1997.

[23] Milind Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124, 1997.

[24] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[25] Michael P. Wellman. Online marketplaces. In M. Singh, editor, *Practical handbook of Internet Computing*. CRC Press, 2004.

[26] Michael P. Wellman, Shih-Fen Cheng, Daniel M. Reeves, and Kevin M. Lochner. Trading agents competing: Performance, progress, and market effectiveness. *IEEE Intelligent Systems*, 18(6):48–53, 2003.

[27] Michael P. Wellman, Jeffrey K. MacKie-Mason, Daniel M. Reeves, and Sowmya Swaminathan. Exploring bidding strategies for market-based scheduling. In *Fourth ACM Conference on Electronic Commerce*, pages 115–124, San Diego, June 2003.

[28] Michael P. Wellman, Daniel M. Reeves, Kevin M. Lochner, Shih-Fen Cheng, and Rahul Suri. Approximate strategic reasoning through hierarchical reduction of large symmetric games. In *Twentieth National Conference on Artificial Intelligence*, pages 502–508, Pittsburgh, 2005.

[29] Michael P. Wellman, Daniel M. Reeves, Kevin M. Lochner, and Rahul Suri. Searching for Walverine 2005. In *IJCAI-05 Workshop on Trading Agent Design and Analysis*, Edinburgh, 2005.

[30] Michael P. Wellman, Daniel M. Reeves, Kevin M. Lochner, and Yevgeniy Vorobeychik. Price prediction in a trading agent competition. *Journal of Artificial Intelligence Research*, 21:19–36, 2004.

[31] Michael P. Wellman, Peter R. Wurman, Kevin O'Malley, Roshan Bangera, Shou-de Lin, Daniel Reeves, and William E. Walsh. Designing the market game for a trading agent competition. *IEEE Internet Computing*, 5(2):43–51, 2001.