

# A Grey-Box Approach to Automated Mechanism Design

(Extended Abstract)

Jinzhong Niu, Kai Cai  
Computer Science  
The Graduate Center, CUNY  
365 Fifth Avenue, New York, NY 10016, USA  
{jniu, kcai}@gc.cuny.edu

Simon Parsons  
Computer and Information Science  
Brooklyn College, CUNY  
2900 Bedford Avenue, Brooklyn, NY 11210, USA  
parsons@sci.brooklyn.cuny.edu

## ABSTRACT

This paper presents an approach to automated mechanism design in the domain of double auctions. We describe a novel parameterized space of double auctions, and then introduce an evolutionary search method that searches this space of parameters. The approach evaluates auction mechanisms using the framework of the TAC Market Design Game and relates the performance of the markets in that game to their constituent parts using reinforcement learning. Experiments show that the strongest mechanisms we found using this approach are able to win the Market Design Game against known, strong opponents.

## Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent systems*

## General Terms

Algorithms, design, economics, experimentation, measurement

## Keywords

Mechanism design, AMD, double auctions, reinforcement learning

## 1. INTRODUCTION

Auctions play an important role in electronic commerce, and have been used to solve problems in distributed computing. A major problem that needs to solve in these fields is: *Given a certain set of restrictions and desired outcomes, how can we design a good, if not optimal, auction mechanism; or when the restrictions and goals alter, how can the current mechanism be improved to handle the new scenario?*

Traditionally, a mechanism is designed by hand, analyzed theoretically, and then revised as necessary. The problems with the approach are exactly those that dog any manual process—it is slow, error-prone, and restricted to just a handful of individuals with the necessary skills and knowledge. In addition, there are classes of commonly used mechanism, such as the double auctions that we discuss here, which are too complex to be analyzed theoretically, at least for interesting cases [14].

Automated mechanism design (AMD) aims to overcome the problems of the manual process by designing auction mechanisms auto-

**Cite as:** A Grey-Box Approach to Automated Mechanism Design (Extended Abstract), Jinzhong Niu, Kai Cai, and Simon Parsons, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX.

Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

matically. AMD considers design to be a search through some space of possible mechanisms. For example, Cliff [2] and Phelps [10, 11] explored the use of evolutionary algorithms to optimize different aspects of the continuous double auction. Around the same time, Conitzer and Sandholm [3] were examining the complexity of building a mechanism that fitted a particular specification.

These different approaches were all problematic. The algorithms that Conitzer and Sandholm considered dealt with exhaustive search, and naturally the complexity was exponential. In contrast, the approaches that Cliff and Phelps pursued were computationally more appealing, but gave no guarantee of success and were only searching tiny sections of the search space for the mechanisms they considered. As a result, one might consider the work of Cliff and Phelps, and indeed the work we describe here, to be what Conitzer and Sandholm [4] call “incremental” mechanism design, where one starts with an existing mechanism and incrementally alters parts of it, aiming to iterate towards an optimal mechanism. Similar work, though work that uses a different approach to searching the space of possible mechanisms has been carried out by [13] and has been applied to several different mechanism design problems [12].

The problem with taking the automated approach to mechanism design further is how to make it scale—though framing it as an incremental process is a good way to look at it, it does not provide much practical guidance about how to proceed.

## 2. GREY-BOX AMD

We propose a *grey-box* AMD approach, which emerged from our previous work on the analyses of the Trading Agent Competition Market Design game, also known as the CAT game [1]. In [8], we examined how the internal design of markets relates to the dynamics of a single game between these markets, and in [6], we viewed markets as atomic entities and compared their performances in various types of CAT match-ups. These two pieces of prior work are respectively analogous to the *white-box* testing and *black-box* testing in software engineering, and have both advantages and disadvantages. The white-box approach is capable of revealing which part of a mechanism may cause vulnerabilities, but it requires internal structure and involves manual examination. The black-box approach does not rely upon the accessibility of the internal design of a mechanism and can be applied to virtually any strategic game. However the black-box approach tells us little about what may have caused a mechanism to perform poorly and provides little in the way of hints as to how to improve the design. The grey-box approach combines the white-box approach and the black-box approach and can automatically create a complex mechanism by searching a structured space of auction components using evolutionary computation and reinforcement learning techniques.

More specifically, we view a market mechanism as a combina-

tion of auction rules, each as an atomic building block, and maintain a population of building blocks. We associate each block with a *quality score*, which reflects the fitnesses of auction mechanisms using this block, explore the part of the space of auction mechanisms that involves building blocks of higher quality, and keep the best mechanisms we find.

The grey-box approach addresses the following issues. First, a set of building blocks for auction mechanisms is a prerequisite for automated mechanism design in the grey-box approach and the quantity and quality of these building blocks to much extent determine how successful the approach may be. To this end, we introduced a parameterized framework of auction mechanisms in the domain of CAT games [8]. The framework includes multiple interwind components, or *policies*, each regulating one aspect of a market. We have collected policies either from the literature [5], or from our previous work [6, 8, 9], or contributed by entrants to the CAT competitions. These policies, each as a building block, can create millions of different combinations, i.e., auction mechanisms, and provide a solid foundation for the grey-box approach.

Second, the parameterized framework and the quality scores enable us to choose better building blocks to create effective auction mechanisms for CAT games. Suppose the set of building blocks is  $\mathbb{B} = \{B_{ij}\}$ , where  $B_{ij}$  is the  $j$ th auction policy about the  $i$ th aspect of auction mechanisms. The group of all the policies on the  $i$ th aspect is denoted as  $B_i$ . Then choosing auction policies to construct one single auction mechanism becomes multiple  $n$ -armed—more accurately,  $|B_i|$ -armed—bandit problems. We used simple methods like  $\epsilon$ -greedy to solve these problems, balancing exploration and exploitation.

Third, CAT games allow us to evaluate an auction mechanism directly against other mechanisms. We sample multiple mechanism from the space at once and put it into a competition against a certain set of fixed markets upon which we desire to improve. We also keep a Hall of Fame that includes the top mechanisms we found in the CAT games over time and put them into the next round of games so as to have moving targets as well.

Finally, the scores of markets in CAT games provide feedback to those building blocks used by the markets. Action-value methods in reinforcement learning—methods that estimate the expected values of actions based on rewards when the actions are selected—come in handy to update the quality scores of the building blocks.

We ran a series of experiments that each lasted 200 steps, used the classic clearing house and continuous double auctions as fixed market competitors, and eventually obtained three strong mechanisms for CAT games. We demonstrated in further experiments that these three mechanisms would have won the 2007 and 2008 CAT competitions if they had participated in the games.

An extended version of this paper can be found in [7].

### 3. SUMMARY

The grey-box approach combines evolutionary computation and reinforcement learning methods and provides a practical guidance to automated mechanism design in a non-trivial setting. It can be used in other strategic, competitive situations where decision making processes are complex and difficult to design and evaluate manually.

### 4. REFERENCES

- [1] K. Cai, E. Gerding, P. McBurney, J. Niu, S. Parsons, and S. Phelps. Overview of CAT: A market design competition. Technical Report ULCS-09-005, Department of Computer Science, University of Liverpool, Liverpool, UK, 2009. Version 2.0.
- [2] D. Cliff. Evolution of market mechanism through a continuous space of auction-types. Technical Report HPL-2001-326, Hewlett-Packard Research Laboratories, Bristol, England, Dec. 2001.
- [3] V. Conitzer and T. Sandholm. Automated mechanism design: Complexity results stemming from the single-agent setting. In *Proceedings of the Fifth International Conference on Electronic Commerce (ICEC'03)*, pages 17–24, Pittsburgh, PA, USA, September 2003. ACM Press.
- [4] V. Conitzer and T. Sandholm. Incremental mechanism design. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, pages 1251–1256, Hyderabad, India, January 2007.
- [5] K. A. McCabe, S. J. Rassenti, and V. L. Smith. Designing a uniform price double auction. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, chapter 11, pages 307–332. Westview Press, Perseus Books Group, Cambridge, MA, 1993.
- [6] J. Niu, K. Cai, P. McBurney, and S. Parsons. An analysis of entries in the First TAC Market Design Competition. In *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT)*, Sydney, Australia, December 2008.
- [7] J. Niu, K. Cai, and S. Parsons. A grey-box approach to automated mechanism design. *Computing Research Repository (CoRR)*, abs/1002.0378, 2010.
- [8] J. Niu, K. Cai, S. Parsons, E. Gerding, and P. McBurney. Characterizing effective auction mechanisms: Insights from the 2007 TAC Mechanism Design Competition. In Padgham, Parkes, Müller, and Parsons, editors, *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, pages 1079–1086, Estoril, Portugal, May 2008.
- [9] J. Niu, K. Cai, S. Parsons, and E. Sklar. Reducing price fluctuation in continuous double auctions through pricing policy and shout improvement rule. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)*, pages 1143–1150, Hakodate, Japan, May 2006.
- [10] S. Phelps, P. McBurney, S. Parsons, and E. Sklar. Co-evolutionary auction mechanism design: a preliminary report. In *Proceedings of the Workshop on Agent Mediated Electronic Commerce IV (AMEC IV)*, 2002.
- [11] S. Phelps, S. Parsons, E. Sklar, and P. McBurney. Using genetic programming to optimise pricing rules for a double auction market. In *Proceedings of the Workshop on Agents for Electronic Commerce*, Pittsburgh, PA, 2003.
- [12] J. L. Schwartzman and M. P. Wellman. Exploring large strategy spaces in empirical game modeling. In *Proceedings of the 11th Workshop on Agent-Mediated Electronic Commerce*, 2009.
- [13] Y. Vorobeychik, D. M. Reeves, and M. P. Wellman. Constrained automated mechanism for infinite games of incomplete information. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, pages 400–407, 2007.
- [14] W. Walsh, R. Das, G. Tesauro, and J. O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In P. Gmytrasiewicz and S. Parsons, editors, *Proceedings of 2002 Workshop on Game-Theoretic and Decision-Theoretic Agents (GTDT-02)*, Edmonton, Alberta Canada, July 2002.