Research Statement

Shih-Fen Cheng
School of Information Systems, Singapore Management University
Tel: (65) 6828-0526; Email: sfcheng@smu.edu.sg
Updated on February 10, 2012

Background

I am academically trained in operations research and artificial intelligence. My research interests in both areas are centered around decentralization, albeit for different reasons. In operations research, the motivation for decentralization is to reduce dimensionality for very large scale problems. In artificial intelligence (multi-agent area in particular), decentralization comes naturally since most problems have decentralized control and distributed information.

Most of my research in recent years has been in agent-based financial markets, urban transportation, and decentralized optimization. In these areas, I enjoy applied research, which includes building systems and working with real-world datasets. My theoretical work is usually done later to explain principles involved in system building and to provide meta-level insights on empirical datasets. Speaking very generally, my primary objectives across all areas are to understand the mechanics of the studied systems and to design and implement novel mechanisms that will improve the performance of these systems. Since each area has its unique issues and backgrounds, I will describe my work in each area separately.

Research Areas

Agent-based Financial Markets

Agent-based approaches in financial markets have been pioneered by researchers at the Santa Fe Institute since early 90s. However, I believe I am among the first few who actually used the technique in training human and analyzing regulatory policies.

My first major contribution in the area is the construction of a software platform that can be generally applied for various purposes. The design principles of this platform have been published in important AI conferences like IAAI [1], AAMAS [2], and IAT [3, 4].

Besides its academic merits, this platform has been used extensively since 2007 within SMU for training both undergraduates and students from the executive programs. This research project was first supported by the International Enterprise Singapore, and it's now continuously supported by the International Trading Institute at SMU.

Besides used in training, the software platform is also utilized for agent-based policy analysis. More specifically, I am interested in investigating the effectiveness of regulatory policies under extreme market conditions. Since most theoretical models break down under extreme conditions, agent-based model is one of the few frameworks that might still work.

One of my recently completed research [5] investigates the potential causes for market instability within commodity markets, and we subsequently test whether the widely debated regulatory measure, limiting speculative position limits, can be an effective policy. The key research findings show that the homogeneity among market participants is what causes extreme price movements, and positions limits cannot contain the induced volatility. The only feasible measure seems to be liquidity injection.

This work was well-received among practitioners, and shortly after its completion, the Institute for Financial Markets (IFM) has awarded us a research grant to analyze the flash crash of May 6, 2010. In this project, we adopt similar generative approach using agent-based technology: we replicated the inter-asset connection and fragmented exchanges, and demonstrated that markets can be forced into unbalances and flash crashes when selling pressure mounts and exchanges are implementing circuit-breakers without coordination [6]. This work has important implications to the regulators everywhere, since we demonstrated that the fragmentation of market exchanges is enough for generating flash crashes, and the inter-connectivity of different asset classes will propagate and amplify the flash crashes from where it starts to a much wider classes of assets.

Future Research

- **Software Infrastructure.** The long-term development goal for the software infrastructure is to allow more complicated and realistic market structure to be simulated. One of the most important such features to be included is the simulation of parallel markets. From the theoretical point of view, parallel markets are no different from a single market where all trades take place. However, in practice, the existence of multiple markets where the trading for an asset can take place introduces information asymmetry, friction, and complicates strategy design. To better understand modern market microstructure, it is thus necessary to have software infrastructure that support this.

- **Agent-based Policy Analysis.** In the near term, my focus in this area will still be on the studying of the potential weakness of modern financial markets. In particular, I will be focusing on replicating popular high-frequency trading (HFT) strategies in detail on our software platform. That will form the foundation for my investigation on analyzing the impact of HFT and also answering questions on how to better manage a marketplace where HFT is the majority.

- **Human Behavior Analysis.** Based on the collected training data over the years, I am working on trying to develop a model for predicting human trader's performance. Compare to similar past researches which are mainly psychological in nature, my research is different in that it focuses on building quantitative models that are solely based on information recorded by the trading platform.

Urban Transportation

My interest in urban transportation started in traffic control, where I looked at the coordinated control of traffic signals in a densely populated urban area. The problem is viewed as a combinatorial optimization problem, and I proposed an efficient distributed algorithm for solving it [7]. More recently, I also begin investigating how the quality of service (QoS) of a taxi fleet can be improved.

Taking taxis is considered by many the public transportation mode that is closet to owning a car since it allows point-to-point movement. Many factors affect the efficiency of taxi service, for example, government policy (quota management in particular), fare structure, and revenue model. In my research, I focus on investigating taxi drivers' reactions to incentives when the fare is distance-based and drivers pay a fixed rent and get to keep all earned revenues.

By analyzing taxi drivers' queueing behaviors at a busy destination, I find out that drivers tend to flock to the queue when the customer traffic is expected to be the highest, and with only local information and no communication, drivers end up suffering the most during the busiest hour (earned less, waited longer) [8]. This study is conducted based on real-world data from a major taxi operator in Singapore and such circumstances are widely observed at high-volume queueing locations with predictable demand patterns. Having characterized the phenomenon, my ongoing effort is on how to resolve it via mechanism design.

I have also constructed an agent-based simulator for taxi fleet operations so that the effectiveness of driver strategies and also the performance of coordination mechanisms can be objectively evaluated [9]. Our simulator, TaxiSim, is the first simulation (to the best of my knowledge) dedicated to the simulation of taxi fleet operation. TaxiSim is capable of supporting the simulation of thousands of taxis simultaneously (each taxi is executed individually as an agent), and it can also incorporate real-world demand data. For most of my follow-on research, TaxiSim will play an important role.

A related study is on the equilibrium strategy in an urban environment. Taxi movement is special compared to ordinary vehicles in that its movement is decided by customers when they are on-board the taxi. The equilibrium of a taxi fleet thus will be dynamic and dependent on the flow of customers [10]. This work provides insight on how to better manage a taxi fleet so that taxis can be re-directed to the areas that need them most, while still observing the autonomy of individual taxis.

Future Research

- **Taxis for the Last-Mile/First-Mile Connection.** In many cities, one critical reason why people don't use public transportation is the lack of means to move to/from the closest access point. To encourage the utilization of public transport, this is one of the most important issues that needs to be addressed. With proper incentives and coordination, taxis can be readily used for this purpose. With the funding support from the Future Urban Mobility Centre of the Singapore-MIT Alliance for Research and Technology, we are actively working on this topic.

- **Coordinated Traffic Control and Routing.** Traffic signals and vehicles are increasing wired by various advanced communication and control technologies. However, little attempt has been made to coordinate the traffic demand, which is affected by vehicle routing, and traffic control. I plan to investigate the use of decentralized approaches in coordinating these two closely related problems in urban environment.

- **Human Behavior Analysis.** The implication and the importance of the taxi problem are well beyond just the transportation domain. Viewed very generally, we are actually studying traces of self-interested individuals in a complex network environment. As already shown in my earlier research, even very experienced human drivers are making inferior decisions and this is probably due to the lack of information and real-time processing capability. My interest in this thread of research is to quantify recurring human fallacies and experiment with approaches that can help us avoiding them.

Decentralized Optimization

Handling large-scale computation and decentralization has been major concerns in all my research areas. To support the investigation in these areas, I also maintain a separate stream of research focusing on developing suitable methodologies and algorithms for these tasks.

There are two major types of methodologies that I am working on:

1. The *market-based approaches* that can deal with problems that are inherently distributed. These problems are usually composed of self-interested individuals (i.e., agents) who possess local information and cannot be controlled centrally. Proper incentives and mechanisms need to be introduced in these cases in order to implement any control policy.

   My methodological contributions in this area are mainly in agent strategy design [11], game-theoretic analytics [12, 13, 14], performance analysis [15], and mechanism design [16, 17].

2. The *game-theoretic algorithms* that can approximately solve high-dimensional hard optimization problem via parallel computation. Roughly speaking, such techniques will

identify manageable sub-problems within a grand problem and regulate how these sub-problems communicate or coordinate with each other. The parallel computing is applied to solve individual sub-problems massively. My interest is to adopt game theory inspired algorithms for such purposes. Example of such approaches include my published work in transportation [7], as well as two projects in production systems that are currently under review [18, 19].

Future Research

The most important issue I am interested in investigating further in this area is how to properly handle uncertainties in distributed environments. Although handling uncertainties has always been an important area of research in classical optimization, how to do it properly in a distributed environment is still not entirely clear. This issue can potentially be addressed by designing novel market mechanisms that deliberately promote robustness. Or it might also be addressed on the agent side, which focus on robust strategy design.

This topic is closely related to my other two interested domains: financial markets and urban mobility. Markets and transportation systems are both decentralized and to improve their performance, we inevitably have to deal with the decentralized optimization problems within them.

Living Analytics

Besides the above three areas that I have worked on for years, I have recently begun looking at how to incorporate social networks into online businesses. The main research question I am interested in is how social network can be better integrated with eCommerce. The first generation of eCommerce takes the form of online catalog, and only very few manage to introduce social elements into it. Amazon is one of few that successfully integrates catalog and social feedbacks in a meaningful manner through collective filtering. More social elements are emerging from the second generation of eCommerce; in fact, many firms have ported the concept in the offline world into the online world. For example, the concepts of auctions, private sales, and collective buyings, which have deep roots in the physical world, have all generated successful online ventures. I believe we are now in the third phase of the eCommerce, in which social networks should eventually play more important role.

The idea of taking advantage of social network is not new in areas like marketing or communication. However, modern technologies make it possible to characterize the whole social network and exert pinpoint control on reaching out to specific influential nodes (e.g., buying up famous bloggers to write ad-blog post). My current interest in this area is to construct representative model that can describe such possibility. The model will be different from the traditional work on influence maximization in that we introduce adversaries who aim to gain influence in the same network. This line of work should shed light on how marketing campaign can be run in the age of social network.

References

[1] Shih-Fen Cheng and Yee Pin Lim. An agent-based commodity trading simulation. In *Twenty-First Annual Conference on Innovative Applications of Artificial Intelligence*, pages 72–78, Pasadena, CA, USA, 2009.

[2] Shih-Fen Cheng, Yee Pin Lim, and Chao-Chi Liu. An agent-based commodity trading simulation (demo paper). In *Eighth International Conference on Autonomous Agents and Multiagent Systems*, pages 1377–1378, Budapest, Hungary, 2009.

[3] Shih-Fen Cheng. Designing the market game for a commodity trading simulation. In *IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, pages 445–449, San Jose, CA, USA, 2007.

[4] Shih-Fen Cheng. Event study method for validating agent-based trading simulations. In *IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, pages 465–469, Toronto, Canada, 2010.

[5] Bernard Lee, Shih-Fen Cheng, and Annie Koh. An analysis of extreme price shocks and illiquidity among systematic trend followers. *Review of Futures Market*, 18(4):385–419, 2010.

[6] Bernard Lee, Shih-Fen Cheng, and Annie Koh. Would price limits have made any difference to the 'Flash Crash' on May 6, 2010. *Review of Futures Market*, 19(Special IFM Issue):55–93, 2011.

[7] Shih-Fen Cheng, Marina A. Epelman, and Robert L. Smith. CoSIGN: A parallel algorithm for coordinated traffic signal control. *IEEE Transactions on Intelligent Transportation Systems*, 7(4):551–564, 2006.

[8] Shih-Fen Cheng and Xin Qu. A service choice model for optimizing taxi service delivery. In *12th International IEEE Conference on Intelligent Transportation Systems*, pages 66–71. St. Louis, MO, 2009.

[9] Shih-Fen Cheng and Thi Duong Nguyen. Taxisim: A multiagent simulation platform for evaluating taxi fleet operations. Technical report, Singapore Management University, 2011.

[10] Pradeep Varakantham, Shih-Fen Cheng, and Thi Duong Nguyen. Decentralized decision support for an agent population in dynamic and uncertain domains (Extended Abstract). In *Tenth International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS-11)*, Taipei, Taiwan, 2011.

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

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

[13] Shih-Fen Cheng and Michael P. Wellman. Iterated weaker-than-weak dominance. In *Twentieth International Joint Conference on Artificial Intelligence*, pages 1233–1238, Hyderabad, India, 2007.

[14] 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, 2005.

[15] 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.

[16] Hoong Chuin Lau, Shih-Fen Cheng, Thin Yin Leong, Jong Han Park, and Zhengyi Zhao. Multi-period combinatorial auction mechanism for distributed resource allocation and scheduling. In *IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, pages 407–411, San Jose, CA, USA, 2007.

[17] Shih-Fen Cheng, John Tajan, and Hoong Chuin Lau. Robust distributed scheduling via time period aggregation. *Web Intelligence and Agent Systems*, 2011.

[18] Archis Ghate, Shih-Fen Cheng, Stephen Baumert, Daniel Reaume, Dushyant Sharma, and Robert L. Smith. Sampled fictitious play for multi-action stochastic dynamic programs. Under revision for 2nd round review at IIE Transactions, 2011.

[19] Shih-Fen Cheng, Blake Nicholson, Marina A. Epelman, Daniel Reaume, and Robert L. Smith. A dynamic programming approach to achieving an optimal end state along a serial production line. Under 2nd round review at IIE Transactions, 2012.