Distributed Coordination and Network Structure: Experimental Results from Simulation

Jamie Olson and Kathleen Carley
Carnegie Mellon University

Abstract

Coordination is an important phenomena occurring in a wide variety of social and technical systems. We use simulation to examine the ways in which one important system property, the interaction network, effects overall levels of coordination. In particular, we survey the performance of six different learning algorithms, including reasonable strategies and no regret strategies on networks generated by six different algorithms. Our results suggest that no-regret mechanisms not only perform better but also come closer to replicating human behavior in the network coordination task.

Introduction

Coordination and cooperation are extremely important social phenomena that have prompted research in a wide array of disciplines, including sociology, psychology, economics and philosophy. Game theorists in particular have been interested in the study of cooperation and coordination. The Prisoner's Dilemma (Axelrod 1984) alone inspired more that one thousand articles (Donninger 1986). More importantly, research on the Prisoner's Dilemma and related problems has been applied to real world scenarios (e.g. (Cable & Shane 1997)). Commons(**?**), a Clearly, these are important problems with significant real-world application.

As awareness of social networking websites like Facebook, MySpace and LiveJournal grows, both researchers and members of the public alike are becoming increasingly aware of the crucial role that social networks play in a wide variety of social phenomena (e.g. (Backstrom *et al.* 2006)). Among other results, there is significant evidence that an individual's position within the social network plays a significant role in influencing their productivity (Cataldo *et al.* 2007).

Parallel to this interest in social network research, a new game theoretic framework has been developed that incorporates some aspects of social networks into a game theoretic framework. These new games, called graphical games, restrict each player's influence to his immediate neighbors in the social network (Kearns 2007). Early experimental results from these games indicate significant network effects. Kearns, Suri, & Montfort (2006) conducted experiments in which human subjects solved a version of the vertex coloring problem that was converted into a graphical game. Each participant had control of a single vertex in the network. One task was to coordinate (indirectly) with their neighbors so that they all had the same color. Participants had limited knowledge of the network; they were only able to see their immediate neighbors. Figure 1 shows an example of this game.

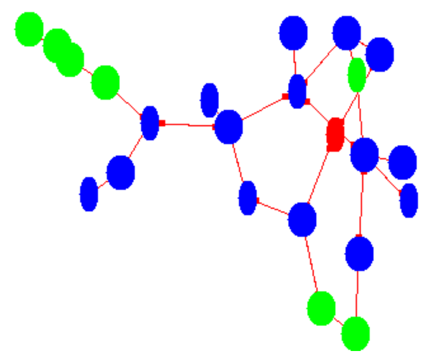

Figure 1: An Example of the Graph Coloring Coordination Problem

Kearns, Suri, & Montfort (2006) observed several interesting relationships between network structure and apparent difficulty of coordination. In particular, they found that networks generated by preferential attachment made solving the coloring problem more difficult than did networks based on cyclical structures, and that small-world networks were easier still. However, due of the constraints imposed by working with human subjects, Kearns, Suri, & Montfort (2006) were able to test a limited of networks. As such, there were several potentially confounding variables, including network size and density.

This work attempts a more robust investigation into the effects of network structure on coordination by using simulation to more completely explore the network dimension.

Background

Learning Mechanisms

We use the modified versions of six learning algorithms as introduced by Greenwald *et al.* (2001). The algorithms can be divided evenly into two classes, reasonable learning methods (Friedman & Shenker 1998; Erev & Roth 1996; Friedman & Shenker 1996) and no-regret learning methods (Foster & Vohra 1993; Auer *et al.* 1995; Hart & Mas-Colell 2000). Reasonable learning algorithms are algorithms that rapidly learn to play (with high probability) the strategy with the highest average payoff (Friedman & Shenker 1998). No-regret learning is based on minimizing a mathematical formalization of regret. In this context, the regret for a player using strategy i is the difference between the payoff actually received for using strategy i and the payoff that would have been received if a different strategy, j, had been used every time that i was used (Greenwald & Jafari 2003).

Network Topologies

A variety of network generation algorithms were used to construct the interaction network. These included two scale-free networks (Barabási & Albert 1999; Eppstein & Wang 2002), two small-world networks (Kleinberg 200; Watts 2003) and two simple networks (lattice and Erdös & Rényi (1960)). A comparison of the network generation algorithms can be found in table 2 [1].

Name	Small-World	Scale-Free
Barabasi-Albert (PA)	Yes	Yes
Eppstein-Wang (Eppstein)	Yes	Yes
Erdos-Renyi (Erdos)	No	No
Kleinberg (Kleinberg)	Yes	No
Watts (Watts)	Yes	No
Lattice (Lattice)	No	No

Figure 2: A Comparison of the network generation models

Methods

The simulation was built using the Repast (North, Collier, & Vos 2006) framework for agent-based simulation. The underlying interaction for the simulation was the coordination version of the graph coloring problem. In this game, agents control the color of a single vertex in a network and are rewarded based on the number of neighbors who have chosen the same color. The underlying network remains unchanged for the duration of the simulation. Each simulation consists of a sequence of 2000 rounds with each round consisting of several stages. First, each agent independently chooses a future action. When all agents have made a decision, they act. In this case, acting consists of revealing the new action and being rewarded based on the new coloring of the network.[2] Finally, after all agents have acted, each agent is given an opportunity to learn from their new observation.

Each simulation consists of a homogeneous population all using the same learning algorithm and utility function [3]

Four independent variables were varied for the simulation: the number of vertices in the network, the average number of neighbors each vertex had(average degree), the graph generation algorithm and the learning algorithm. The numbers of vertices in the network varied between 49, 100, 400 and 900, while the average degree was one of 2.0, 4.0, 8.0, 16.0 and 24.0.

Results

The results provide strong evidence that the structure of the interaction network can influence coordination. In particular, there is evidence that lattice-style networks lead to lower levels of coordination. In addition, no-regret algorithms in general and the Foster-Vohra and Internal Regret algorithms in particular show both .

Learning Algorithms

We begin by comparing the performance of the assorted learning methods. Figure 3 shows the performance of each learning method. It would appear that the no regret algorithms and potentially the stage learning algorithm are doing the best.[4] Indeed, when the performance of each learning method is separated by graph topology we can see that the FV and IR algorithms outperform the others by roughly a factor of 10 (see Figure 4). This is again the case when examining the robustness of the learning algorithms as the graph increase in size (Figure 5) and density (Figure 6). Table 1 shows the impact as determined by a simple linear model of each learning method compared to random. Although all algorithms do perform significantly better than random, the difference is not compelling outside of the FV and IR.

Learning Method	Impact
RE	108
Responsive	112
Stage	264
FS	610
FV	1399
IR	1435

Table 1: Average number of conflicts compared to random.

[1] The network generation parameters used were $\beta = 0.5$(Watts) and $\alpha = 2.0$(Kleinberg).

[2] Agent behavior was synchronized for simplicity. Future work includes removing this simplification.

[3] The reward function is independent of the "color" chosen. In other words there is no inherent preference for one color over other. The informed, responsive variants of the algorithms were used with parameters $\gamma = 0.45$, $\epsilon = 0.05$, $\kappa = 5$(IR), $\beta = 1.0$(FS), and $\alpha = 1000.0$(FV). These were the best-performing parameters for each algorithm.

[4] Because of the number of simulations run, all differences mentioned are statistically significant($p < 0.01$).

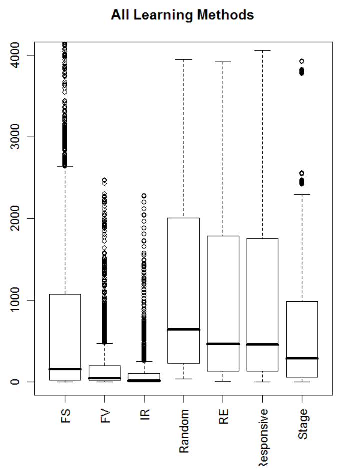

Figure 3: Average Number of conflicts as a function of the learning method.

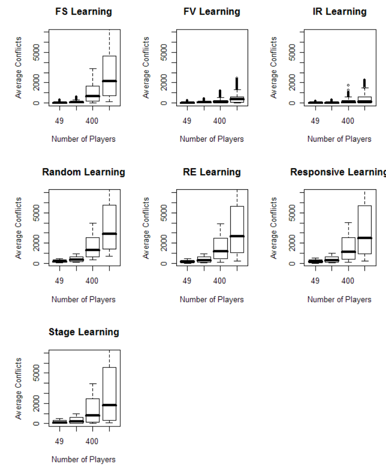

Figure 5: A sequence of boxplots illustrating the influence that the size of the network had on the average number of conflicts. Each plot shows data for a single learning algorithm.

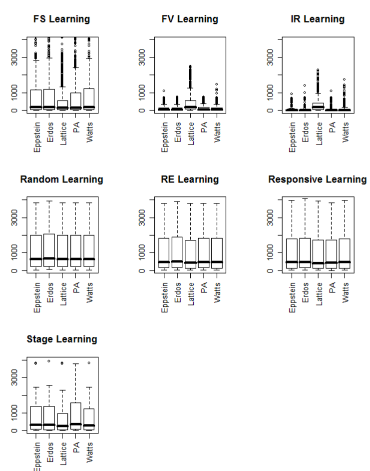

Figure 4: A sequence of boxplots illustrating the influence that the underlying network generation algorithm had on the average number of conflicts. Each plot shows data for a single learning algorithm.

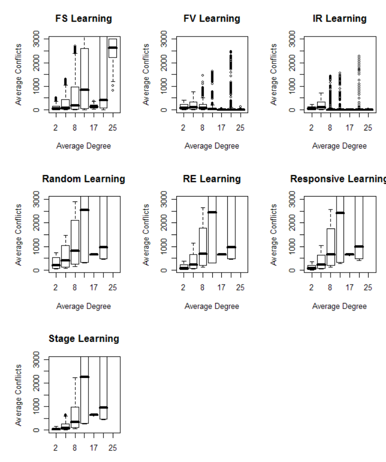

Figure 6: A sequence of boxplots illustrating the influence that the density of the network had on the average number of conflicts. Each plot shows data for a single learning algorithm.

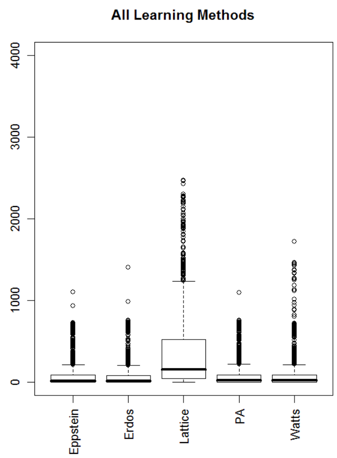

Figure 7: Performance of the FV and IR learning methods as a function of network topology.

Network Topology

The effect of network topology appears to be highly dependent on the learning method being used. However, because FV and IR perform much better and more consistently, we focus on these learning methods. Figure 7 shows the performance of FV and IR on the assorted networks. As you can see, performance is severely degraded on the lattice network topology. This mirrors the performance of human subjects in this task and may lend ecological validity to complement the performance advantage of no-regret learning methods.

Conclusion & Discussion

Interaction networks do indeed have an effect on overall coordination levels. However, the precise effect depends on the learning strategy. With no-regret learning methods, lattice-style networks tend toward significantly lower levels of coordination than other types of networks. Surprisingly, this effect only holds for lattice networks and not for non-small-world networks in general(e.g. Erdös-Rényi networks).

Further work is needed in relaxing some of the assumptions being made. In particular, the assumption regarding synchronous decision-making among the population is clearly unrealistic and must be removed. Also, the assumption of a homogeneous population, regarding both learning algorithms and preferences among the various “colors” must be relaxed. In addition, it is unclear exactly why the reasonable strategies performed so poorly. More experiments are needed to determine precisely the set of circumstances in which reasonable strategies underperform. Despite these shortcomings, this work provides additional evidence that no-regret algorithms consistently perform well in a network context. In addition, there is evidence that no-regret algorithm are capable of mimicking observable human behavior.

References

Auer, P.; Cesa-Bianchi, N.; Freund, Y.; and Schapire, R. 1995. Gambling in a rigged casino: The adversarial multi-armed bandit problem. *focs* 00:322.

Axelrod, R. 1984. *The Evolution of cooperation*. New York: Basic Books, Inc.

Backstrom, L.; Huttenlocher, D.; Kleinberg, J.; and Lan, X. 2006. Group formation in large social networks: membership, growth, and evolution. In *Proceedings of the 12th International Conference on Knowledge Discovery and Data Mining*, 44–54. New York, NY, USA: ACM.

Barabási, A.-L., and Albert, R. 1999. Emergence of scaling in random networks. *Science* 286(5439):509–512.

Cable, D., and Shane, S. 1997. A prisoner's dilemma approach to entrepreneur-venture capitalist relationships. *The Academy of Management Review* 22:142–176.

Cataldo, M.; Bass, M.; Herbsleb, J. D.; and Bass, L. 2007. On coordination mechanism in global software development. *Proceedings of the International Conference on Global Software Engineering*.

Donninger, C. 1986. Is it always efficient to be nice? In Diekmann, A., and Mitter, P., eds., *Paradoxical Effects of Social Behavior: Essays in Honor of Anatol Rapoport*. Heidelberg: Physica Verlag. 123–134.

Eppstein, D., and Wang, J. Y. 2002. A steady state model for graph power laws. In *2nd Int. Worksh. Web Dynamics*.

Erdös, P., and Rényi, A. 1960. The evolution of random graphs. *Magyar Tud. Akad. Mat. Kutat Int. Kzl.* 17–61.

Erev, I., and Roth, A. 1996. On the need for low rationality cognitive game theory: Reinforcement learning in experimental games with unique mixed strategy equilibria. Technical report.

Foster, D. P., and Vohra, R. V. 1993. A randomization rule for selecting forecasts. *Operations Research* 41(4):704–709.

Friedman, E., and Shenker, S. 1996. Synchronous and asynchronous learning by responsive learning automata.

Friedman, E., and Shenker, S. 1998. Learning and implementation on the internet. Departmental Working Papers 199821, Rutgers University, Department of Economics. available at http://ideas.repec.org/p/rut/rutres/199821.html.

Greenwald, A., and Jafari, A. 2003. A general class of no-regret learning algorithms and game-theoretic equilibria.

Greenwald, A.; Friedman, E.; ; and Shenker, S. 2001. Learning in network contexts: Results from experimental simulations. *Games and Economic Behavior: Special Issue on Economics and Artificial Intelligence* 35:80123.

Hart, S., and Mas-Colell, A. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68(5):1127–1150.

Kearns, M.; Suri, S.; and Montfort, N. 2006. An experimental study of the coloring problem on human subject networks. *Science* 313:824–827.

Kearns, M. 2007. *Algorithmic Game Theory*. Cambridge University Press. chapter Graphical Games.

Kleinberg, J. M. 200. Navigation in a small world. *Nature* 406:845.

North, M. J.; Collier, N.; and Vos, J. 2006. Experiences creating three implementations of the repast agent modeling toolkit. *ACM Transactions on Modeling and Computer Simulation* 16(1):1–25.

Watts, D. J. 2003. *Small Worlds : The Dynamics of Networks between Order and Randomness (Princeton Studies in Complexity)*. Princeton University Press.