An Examination of a Firm's Adaptive Behavior and its Network Structure

Joseph K. DeRosa
MITRE Corporation, Bedford, MA USA
jderosa@mitre.org

Ben Mazzotta
Fletcher School of Law and Diplomacy, Tufts University, Medford, MA USA
benjamin.mazzotta@tufts.edu

Joseph T. Lizier
School of Information Technologies, The University of Sydney, NSW, Australia
CSIRO Information and Communications Technology Centre, North Ryde, NSW Australia
jlizier@it.usyd.edu.au

Luciano Oviedo
Intel Corporation, Rio Rancho, NM USA
luciano.c.oviedo@intel.com

Alexander A. Shpunt
Massachusetts Institute of Technology, Cambridge, MA USA
ashpunt@mit.edu

Carver Tate
Georgetown University, Washington, DC, USA
carvertate@gmail.com

Abstract: The relationship between a generalized firm structure and its ability to adapt to different internal and external influences is examined. We use an information sequence to represent a corporate goal, some innovation within the firm, or a widget being produced. The flow of this information is then measured as a function of different network structures, employee decision making algorithms, and where the information originates. As a result, we gain insight into the interplay between a firm's behavior and network structure in the achievement or propagation of its objectives.

FINAL REPORT

Santa Fe Institute
Complex Systems Summer School
Santa Fe, NM
June 2007

Table of Contents

Introduction

The notion of Innovation has many definitions. John Kao defines "Innovation is creativity applied with intention, to create value. Creativity is the ability of individuals, teams and organizations to generate new ideas, approaches or concepts." [7]

Our model examines the dynamics of information flow through an organization as a function of its structure and the way in which workers in the organization modify their characteristics as a result of being exposed to the information. When that information represents innovation, the model provides a measure of the dynamics of value creation in the firm. We draw on Bednar and Page's exploration [1] of cultural change in a network through interactions among agents. In their model, agents continually alter their state vector of cultural values, trying both to conform to the values of others, and having their values remain internally consistent. Our agents, rather than mixing at random, communicate with their peers, supervisors, and subordinates through an organizational structure that represents a firm. We record the dynamics of information flow around this firm, varying the structure of the firm's hierarchy and cross-department/division links, and also varying the cultural parameters that govern change in the firm.

Model

The firm is composed of actors in a hierarchic structure, modeled using NetLogo [8]. A screenshot from the model is shown in Figure 1. Each actor possesses a vector of state variables. While we referred to this vector as the *widget vector* in our discussions, we discussed several ways to interpret the vector, and a few key generalizations of the widget vector that may make it more useful. The widget vector could represent the product that the business unit makes. It could represent the different activities that a team member must undertake to participate in the running of the business. It could reflect a set of core values, or corporate culture, that firm members must agree on.

The state variables are natural numbers (i.e., non-negative integers). The smallest domain we investigated is (0, 1), and the largest was (0, 10). Similarly, the number of state variables was set between 1 and 10. Both of these assumptions reduce computational complexity, but are fundamentally arbitrary.

The structure of the firm was governed by three variables, *width depth* and *randomization*. The width of a network is the number of direct reports to any given supervisor. The depth refers to the number of layers of hierarchy in any given chain of authority. Each agent is connected to their supervisor, direct-reports and their peers (who are connected to the same supervisor). The randomization refers to the probability of breaking peer links (or edges, in the terminology of graph theory) to neighbors at random and replacing those with links to other on the corporate hierarchy, chosen at random and not necessarily of the same rank. Again, these simplifying assumptions are highly unrealistic, but permitted computation.

At each time step, agents had four actions (decisions or options) available to them for one of their state variables (randomly chosen):

1. match the value of the state variable to the supervisor's value (obey);
2. match the value of the state variable to a peer's value (conform);
3. match the value of the state variable to a direct-report's value (listen);
4. do nothing (no action).[1]

This decision produced three free parameters on the system. The probability of matching a supervisor is denoted by p_s. The probability of matching a peer is denoted by p_p. The probability of matching a direct-report is denoted by p_d. The probability of doing nothing is p_n where

(1) $$p_n = 1 - p_s - p_p - p_d.$$

Later versions of the model changed the structure of the network hierarchy and the interaction rules. In one, we allowed the agents to choose to increase the consistency of their internal variables rather than match with a neighbor. In another, we froze one actor's state vector and waited for the others to conform to it. In still another, we assigned weights to the links to alter the probability that agents would interact with neighbor on those links. These links weights were adaptive, rewarding agents for "good" performance related to an external goal imposed on the firm.

Method

In our modeling, we randomized the initial values of the state vector for the agents, and then recorded how long the system took to converge to a single state vector value for all agents in the firm.

Time to convergence is measured in NetLogo modeling time steps. NetLogo has a clock that records time in discrete steps. Each agent acts, in our model, once at each step of the clock.

For each test described below, we identified ranges of the relevant parameter values. At each combination of parameter values, we ran ten trials of the model and recorded the total times to convergence.

Test 1: Simple Hierarchy

In the first test, we examined the propagation of information through a simple and static hierarchy without network randomization.

We ran ten trials of each combination of widths (2, 3, 5, 8) and depths (2, 3, 5). For these trials, $p_s = 0.5$, $p_p = 0.25$, and $p_d = 0.25$. The possibility of no action was set to zero.

---

[1] We had hoped to include a fifth alternative: change the value of a state variable to a random value. None of the models with this alternative produced interesting results.

We ran ten trials of each value of p_s (0.1, 0.2, 0.3, 0.4, 0.5), p_p (0.1, 0.2, 0.3, 0.4, 0.5), and p_d (0.1, 0.2, 0.3, 0.4, 0.5). For these tests, the default values of the parameters were width = 5, depth =5, p_s = 0.5, p_p = 0.25, and p_d = 0.25. Again, the possibility of no action was set to zero. The sum of the probabilities had to equal zero, so some offsetting adjustments were required on each run.

In this test, the state vector contained a single variable with values between 0 and 9.

Test 2: Random Edges

In the second test, we introduced a new action for each agent: internal consistency. This action was selected with probability p_c, before the external interactions were selected as before, giving:

$$p_n = 1 - p_c - (1 - p_c)(p_s + p_p + p_d). \tag{2}$$

When the agent randomly chose to execute the internal consistency action (as per [1]), they would set the chosen state variable to match the value of another internal state variable chosen at random. The availability of this action sets up a competing dynamic between "desires" for matching neighbors and internal consistency. As such, this experiment is an attempt to extend the "conformity-consistency" experiment with full interaction in [1] onto a network model.

The fraction of links replaced by random links ranged from 0 to 1 in increments of 0.1. We also ran a similar trial for a fully connected graph, i.e., one where all possible edges were drawn in, to attempt to replicate the results for full interaction in [1].

For each of these firms, we ran ten iterations of the test for each a range of internal consistency probabilities p_c between 0 and 0.9, recording the time to convergence as the dependent variable.

In this test, the hierarchy had a depth and width of 4, and the state vector contained ten variables with values between 0 and 5. If internal consistency was not selected, the agent would continue on to select their action using the following probabilities: p_s = 0.13, p_p = 0.44, and p_d = 0.44. The probability of doing nothing was set to zero.

Test 3: Privileged Information Source

For this test, we fixed the state vector value for one layer of agents in the firm. The idea was to represent the phenomenon in a firm where there was a tendency to have privileged or more accurate information; for example, a sales rep interacting with customers and learning what widget would sell better. This could also represent a certain segment of the firm having a propensity for influence, decision making, innovation or goal setting. We ran the model until the whole firm finally agreed on that value for the state vector.

In this test, we used 10% random links in the hierarchy, and a network of depth 5 and width 5. The state vector contained three variables with values between 0 and 1. The layer selected for insertion of the fixed state vector was varied between levels 1 (to represent “top-down” as the source), 2 (to represent “middle-out” as the source) and 4 (to represent “bottom-up” as the source).

Test 4: Adaptive Edge Weights

Agents in this test incorporated a type of learning, which was a radical departure from the earlier zero-intelligence models. At each time step, their behavior was adapted depending on whether their previous action helped bring them closer to the firm’s goal. Agents did not know the value of the goal, but their probabilities for undertaking future actions were adapted depending on this goal.

This was implemented by varying the probabilities used by each agent to select whether to obey to their supervisor, conform to a peer or listen to direct-reports. If adapting to a neighbor resulted in a match to the corresponding variable in the goal vector, the probability of interacting with that neighbor’s group (i.e. supervisor, peers or direct-reports) was increased by a positive reinforcement factor a_p. Alternatively, if the adaptation to a neighbor did not result in a match to the corresponding variable in the goal vector, the probability of interacting with that neighbor’s group (i.e. supervisor, peers or direct-reports) was reduced by a negative reinforcement factor a_n. In the future, more highly weighted links would be visited more often than low-weighted links.

A real world example here is one employee conforming to certain behaviors of a peer: the employee may subsequently be found to be performing better or worse and receive certain rewards or otherwise. They may not know which specific behaviors produced the change, but certainly they know it was related to conforming to their peer. While the model implemented here is again simplistic, it does capture aspects of real-world corporate adaptation.

In this experiment, we used a constant hierarchy with a depth and width of 4, and the state vector contained five variables with values between 0 and 5. We set *initial* values for $p_s = 0.13$, $p_p = 0.44$, and $p_d = 0.44$. For each agent, these values were multiplied by the appropriate reinforcement factor a_p for positive reinforcement or the same factor a_p with a multiplier m_n for negative reinforcement at each time step.

Results

For the most part the results were unsurprising. We were able to identify chaotic behavior in one region of the test of adaptive behavior.

In a simple hierarchy, the fastest way to achieve conformity is to propagate a signal from the top node directly down through all the other layers of the firm (as per Figure 2 and Figure 3). Figure 4 shows that increasing the depth of the organization has a much greater effect in increasing the convergence time than increasing the width: this is because increasing the depth not only increases the number of employees in the organization faster but also introduces greater segregation between employees.

With the introduction of the internal consistency rule, we observed similar behavior described in [1]. Where Bednar et al expected to find the familiar U shape from [1] in a fully connected graph, they actually found increasing behavior of convergence time with increasing p_c until p_c approaches 1. In attempting to replicate this (see Figure 6), we found a similar shape but did not reproduce the significant increase in convergence time with p_c; it is unclear why this was the case. With a hierarchical interaction structure (Figure 5Figure 54 and Figure 6), we observed some similarity in the shape of the results from the full interaction model, though there were several important variations. The convergence times were longer and more uncertain on the hierarchical structure, particularly at low network randomness. Interestingly, the hierarchy without random connections showed the most similarity to the U shape result for a fully connected network in [1], while the flattened U without a peak at large p_c found for maximum random connections on the graph was most similar to our own result for a fully connected network. The latter result is as somewhat as expected since maximal network randomness added to the hierarchy gives rise to small-world effects [2]-[3] where average path length of the network (important for disseminating information) is equivalent to that for a randomly connected network. While this is subtly different from a fully connected network, the larger path length is likely to be offset by larger stability from interacting with a smaller number of colleagues.

The privileged information experiment showed a mildly counterintuitive result (see Figure 7Figure 76). Whereas obedience was the most important parameter of the system in earlier tests, the importance of the top level of the hierarchy was greatly diminished in this test. The top level of the hierarchy was a single agent, whereas the bottom contained hundreds of agents. Fixing hundreds of agents (instead of a single agent) greatly compressed the time to convergence, swamping the effect of obedience with all but the highest values of the obedience parameter. It would be interesting to revisit this experiment using the same number of privileged information sources at each level so as to control for this effect.

An interesting result occurred in test 4, which explored weighted edges. Rather than retaining fixed communication links in the firm, agents could strengthen or weaken links based on their recent performance. The results are displayed in Figure 8Figure 87. Overall, we note a phase change (e.g. see [4],[5], and [6]) where the convergence time was lower for *intermediate* positive reinforcement and negative reinforcement, with the interplay between the two determining the intermediate values required of one variable given the other. It appears that using too high a value for either variable results in divergence in the hierarchy. With highly negative reinforcement of edge weights for example, chaotic dynamics ensued, because a significant proportion of employees stopped interacting with certain neighbors leading to isolation of pockets of the corporation. Without effective communication, the corporation fails to achieve convergence in a reasonable period of time. The huge volatility in mean convergence times for the largest negative reinforcement multipliers (m_n = 1.0 in particular, and for multipliers 0.5, 0.6 and 0.7 at larger values of a_p) reflects both a chaotic dynamic, with sensitive dependence on initial conditions, and also an artifact of the computation. The trials were not allowed to run indefinitely, so where the mean convergence times appear to dip back to lower values for these multipliers, the times recorded may not reflect the true time that would be required. Notice that to avoid divergence, the negative reinforcement had to be weaker than the positive reinforcement (using $m_n < 1$), since at early times in the simulation an employee is significantly more likely to give information that doesn't match the goal vector. Finally, it would be useful to repeat this experiment where the reinforcement was applied to each individual link, rather than to the set that the link was associated with (e.g. the set of peers).

Conclusions

The original concept for the models was to construct a zero-intelligence model of the firm, to place that firm in a competitive environment, and to allow the firm to evolve a structure that suited it. As the work progressed, we experimented with different methods to introduce intelligence into the firm in a limited way.

This constitutes an alternative to traditional models of full information and zero-intelligence models. These first steps are crude at best for controlling the introduction of information into the firm, but appear to be worth exploring in a more rigorous way.

We found that it is possible to force virtually any structure to converge very quickly on some idea with sufficient obedience. The parameters on obedience and listening to subordinates essentially control the rate at which information flows vertically through an organization.

When including the action of internal consistency, the network structure itself does not fundamentally alter the Bednar-Page result described in the introduction, although some interesting subtle differences are observed.

Our investigation of the effect of hierarchical level in placing privileged information sources in the firm was dominated by the effect of level size – this experiment should be repeated using a fixed number of privileged sources at the level of insertion.

Adaptive edge weights behave in largely unsurprising ways, including a phase change with the possibility of splintering the firm with large negative reinforcement weights in interaction with larger positive reinforcement. If negative weights dominate the structure of the firm, pockets can develop without connection to the rest of the firm. Once these pockets have formed, there is a real possibility that the firm will simply never converge to a common widget. It would be interesting to revisit this experiment and adapt edge weights on a per link basis rather than per set (i.e. currently we penalize or reward the whole set of peers or direct-reports based on the information received from one of them).

Acknowledgements

This work was partially supported by the Santa Fe Institute through NSF Grant No. 0200500 entitled "A Broad Research Program in the Sciences of Complexity."

JL acknowledges travel support from the ARC Complex Open Systems Research Network (COSNet) and the CSIRO Complex Systems Science Theme.

Figures

Figure 1. NetLogo model of the Enterprise hierarchical structure

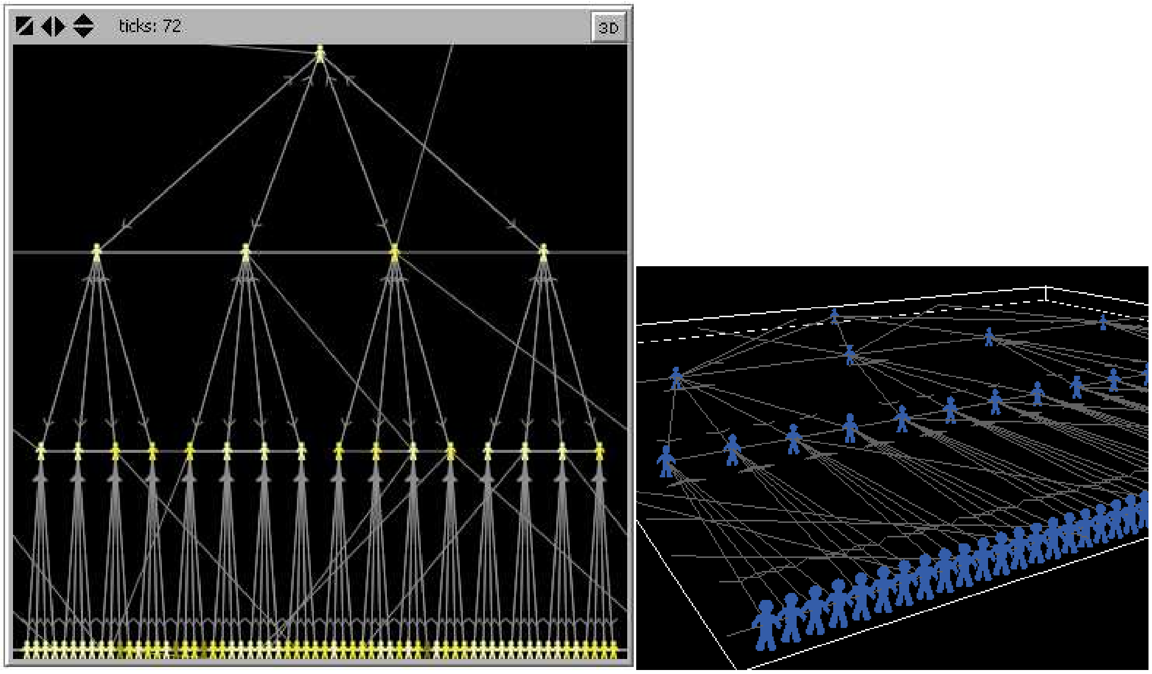

Figure 2. Convergence Time for a Simple Hierarchy

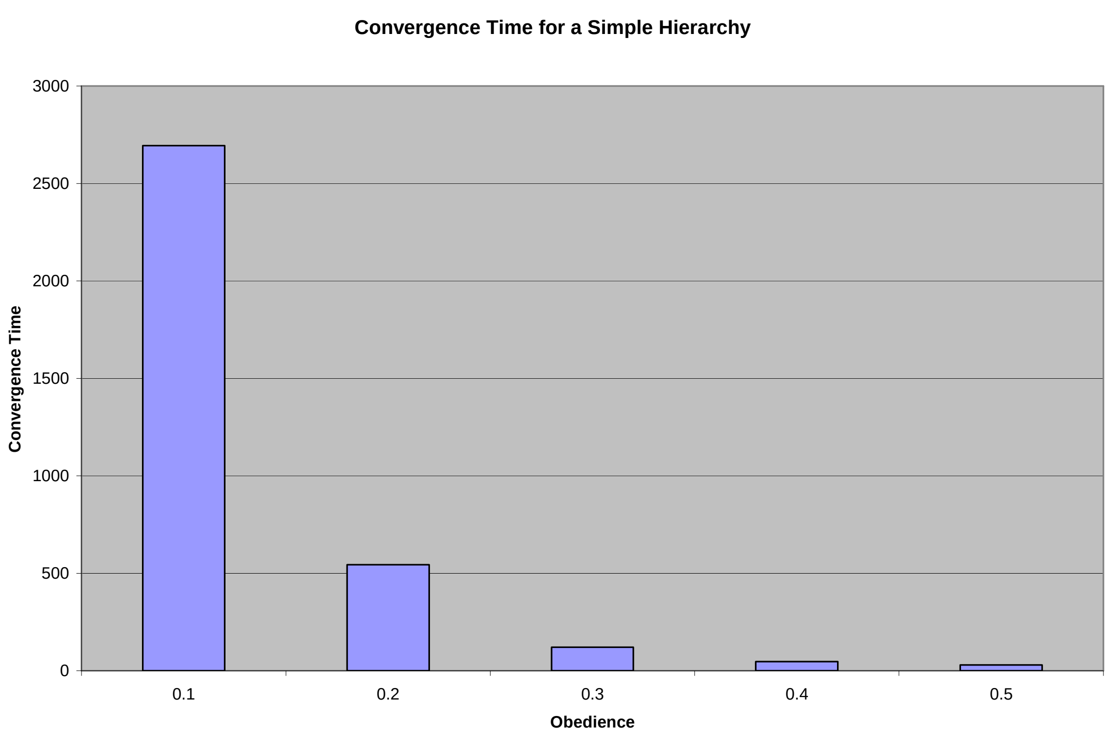

Figure 3. Convergence Time and Conformity

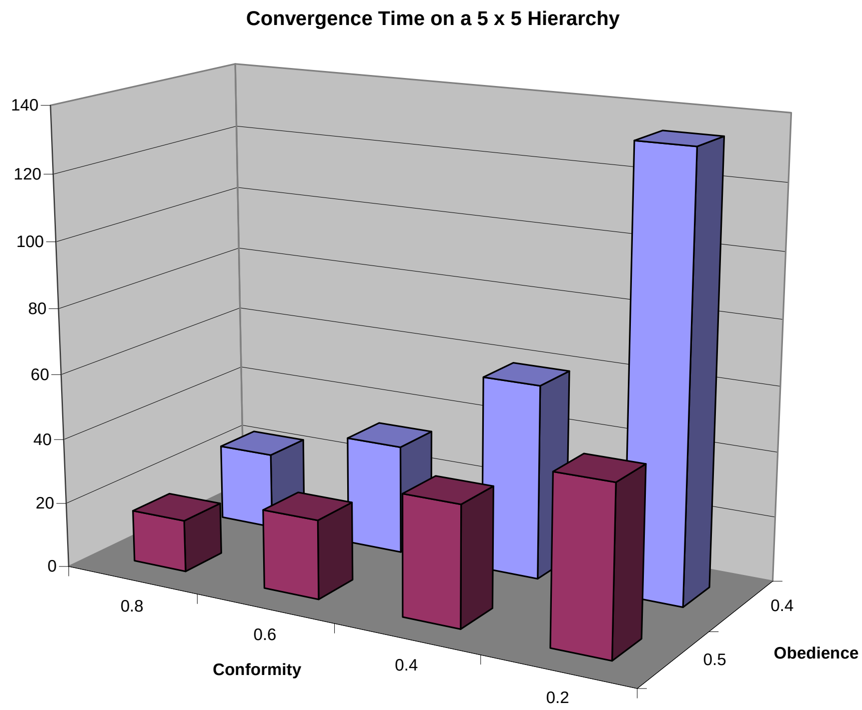

Figure 4. Conformity and Firm Structure

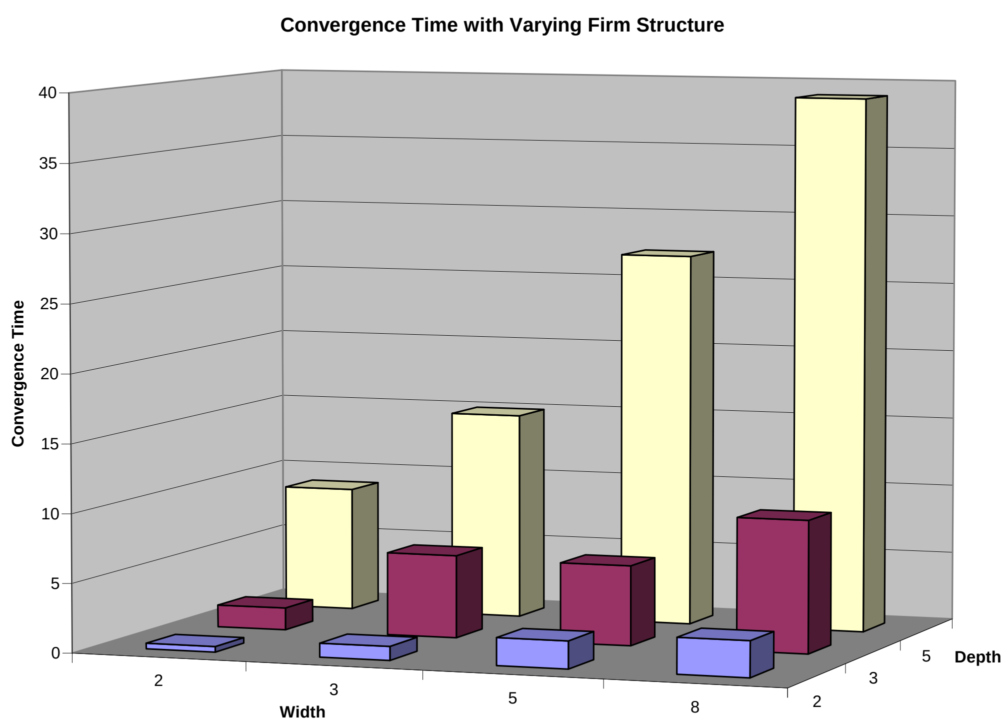

Figure 5. Convergence Time with Internal Consistency rule

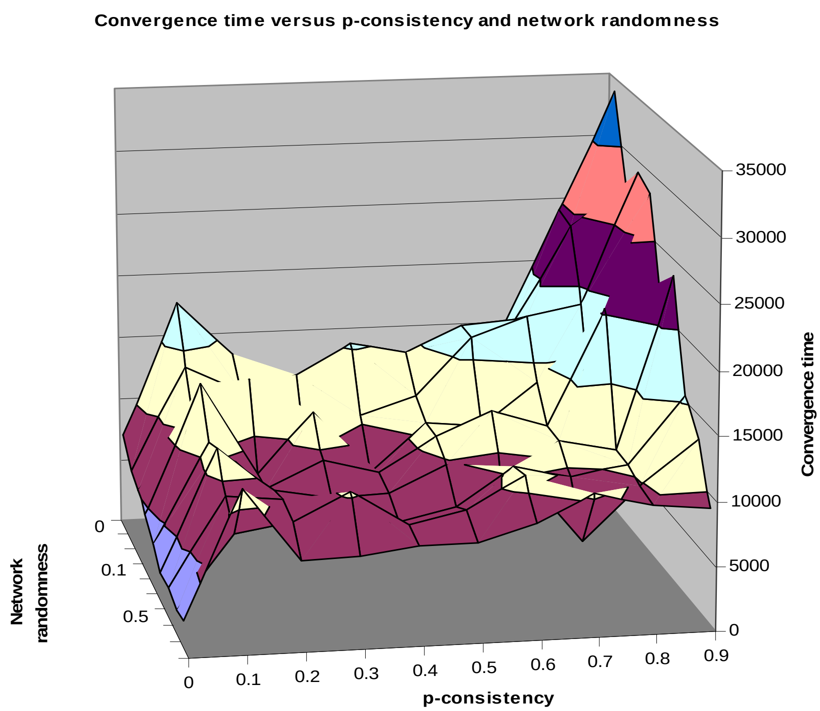

Figure 6. Convergence Time for an Internal Consistency rule, including fully connected graph

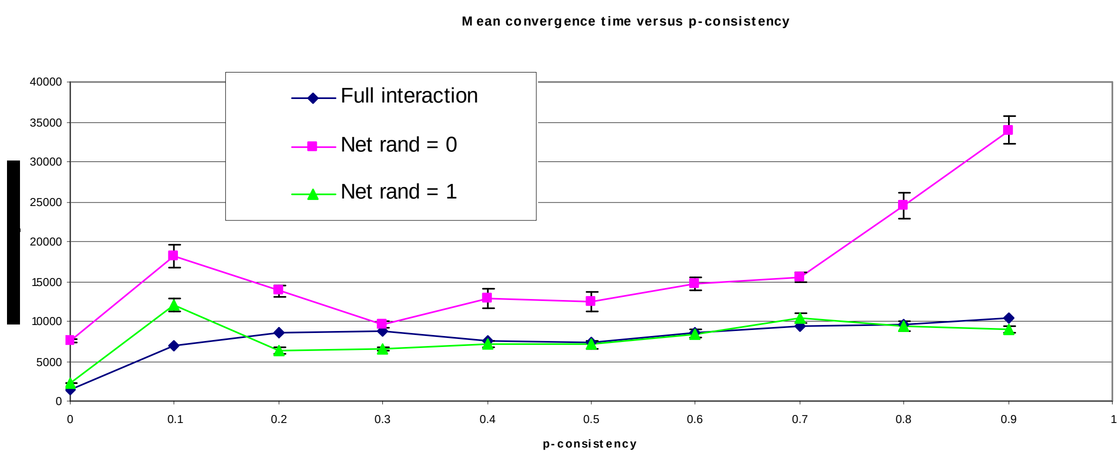

Figure 7. Convergence Time with a Privileged Information Source

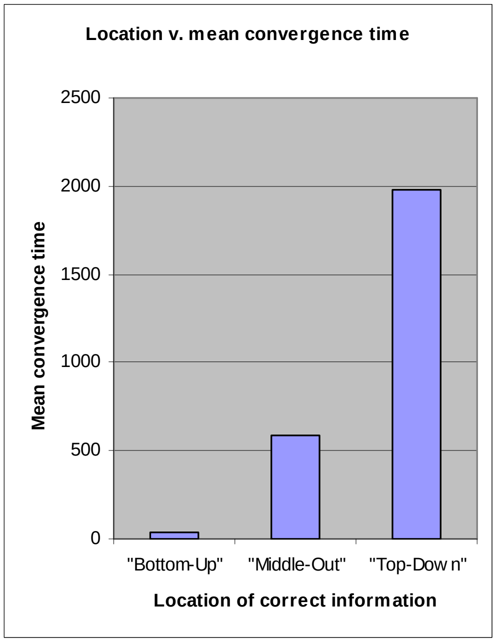

Figure 8. Convergence Time with Adaptive Edge Weights

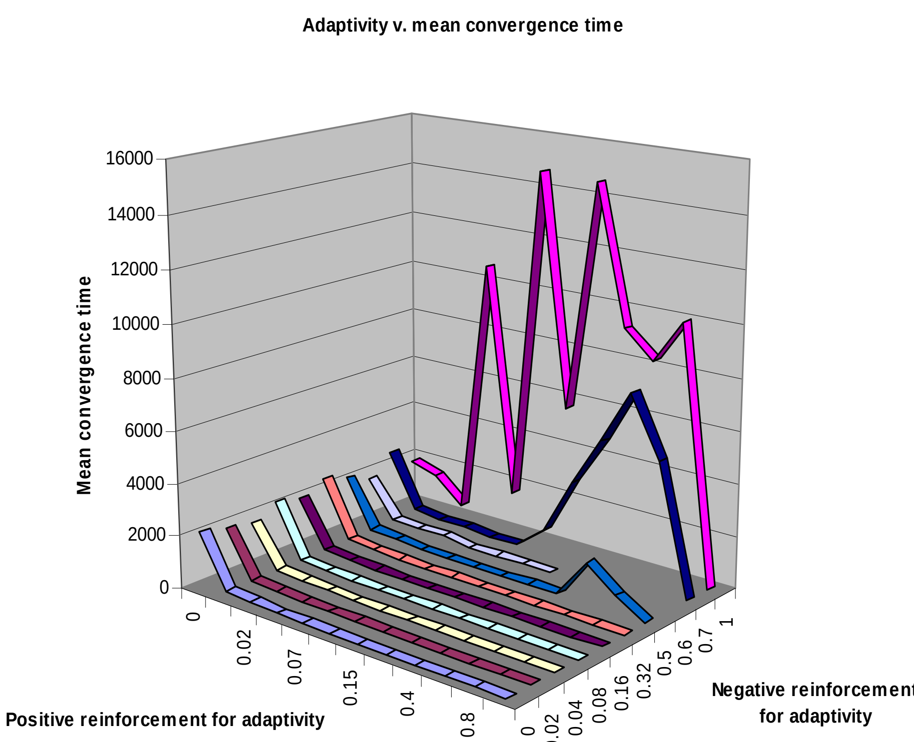

Reference List

[1] J. Bednar, A. Bramson, A. Jones-Rooy, and S. Page, “Consistency, Conformity, and Cultural Heterogeneity”, The University of Michigan, 2006.

[2] D. J. Watts and S. Strogatz, “Collective Dynamics of ‘Small-World’ Networks”, Nature, vol. 393, pp. 440-442, 1998.

[3] D. J. Watts, “Six Degrees: The Science of a Connected Age”. New York: Norton, 2003.

[4] O. Miramontes, “Order-disorder transitions in the behavior of ant societies”, Complexity, vol. 1, pp. 56-60, 1995.

[5] A. Wuensche, “Classifying cellular automata automatically: Finding gliders, filtering, and relating space-time patterns, attractor basins, and the Z parameter”, Complexity, vol. 4, pp. 47-66, 1999.

[6] M. Prokopenko, P. Wang, and D. Price, “Complexity Metrics for Self-monitoring Impact Sensing Networks”, presented at 2005 NASA/DoD Conference on Evolvable Hardware (EH'05), 2005.

[7] J. Kao, Innovation Manifesto, San Francisco: GBN, 2002

[8] U. Wilensky, “NetLogo”, Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, IL, 1999, http://ccl.northwestern.edu/netlogo/