

Self-Organized Criticality of Belief Propagation in Large Heterogeneous Teams

Robin Grinton, Praveen Paruchuri, Paul Scerri, Katia Sycara

Robotics Institute, Carnegie Mellon University**

Summary. Large, heterogeneous teams will often be faced with situations where there is a large volume of incoming, conflicting data about some important fact. Not every team member will have access to the same data and team members will be influenced most by the teammates with whom they communicate directly. In this paper, we use an abstract model to investigate the dynamics and emergent behaviors of a large team trying to decide whether some fact is true. Simulation results show that the belief dynamics of a large team have the properties of a Self-Organizing Critical system. A key property of such systems is that they regularly enter critical states, where one additional input can cause dramatic, system wide changes. In the belief sharing case, this criticality corresponds to a situation where one additional sensor input causes many agents to change their beliefs. This can include the entire team coming to a "wrong" conclusion despite the majority of the evidence suggesting the right conclusion. Self-organizing criticality is not dependent on carefully tuned parameters, hence the observed phenomena are likely to occur in the real world.

1.1 Introduction

The effective sharing and use of uncertain information is key to the success of large heterogeneous teams in complex environments. Typically, noisy information is collected by some portion of the team and shared via the social and/or physical networks connecting members of the team [Bellur *et al.*2002]. Each team member will use incoming, uncertain information and the beliefs of those around them to develop their own beliefs about relevant facts. For example, in the initial reaction to a disaster, members of the response team will form beliefs about the size, scope and key features of the disaster. Those beliefs will be influenced both by what they sense in the environment and what they are told by others in the team. Human teams have developed processes and strategies for managing this messy information and acting without necessarily having convergent beliefs [Sukthankar *et al.*2008]. However, in heterogeneous teams, where team members are not exclusively humans, but may be

** This research has been sponsored by AFOSR grants FA95500810356 and FA95500710039

intelligent agents or robots, and novel network structures connect the team members, we cannot assume that the same tactics will help convergence or that undesirable and unexpected effects will not be observed. Thus, before such teams are deployed in important domains, it is paramount to understand and potentially mitigate any system-wide phenomena that affect convergence. There have been previous attempts in the scientific literature to describe the information dynamics of complex systems, however, due to the complexity of the phenomena involved, mathematical formulations have not been expressive or general enough to capture the important emergent phenomena.

To investigate the dynamics and emergent phenomena of belief propagation in large heterogeneous teams, we developed an abstracted model and simulator of the process. In the model, the team is connected via a network with some team members having direct access to sensors and others relying solely on neighbors in the network to inform their beliefs. Each agent uses Bayesian reasoning over beliefs of direct neighbors and sensor data to maintain a belief about a single fact which can be *true* or *false*. The level of abstraction of the model allows for investigation of team level phenomena decoupled from the noise of high fidelity models or the real-world, allowing for repeatability and systematic varying of parameters. Simulation results show that the number of agents coming to the correct conclusion about a fact and the speed of their convergence to this belief, varies dramatically depending on factors including network structure and density and conditional probabilities on neighbor's information. Moreover, it is sometimes the case that significant portions of the team come to have either no strong belief or the wrong belief despite overwhelming sensor data to the contrary. This is due to the occasional reinforcement of a small amount of incorrect sensor data from neighbors, *echoing* until correct information is ignored.

More generally, the simulation results indicate that the belief propagation model falls into a class of systems known as Self Organizing Critical (SOC) systems [Bak *et al.* 1983]. Such systems naturally move to states where a single local additional action can have a large system wide effect. In the belief propagation case, a single additional piece of sensor data can cause many agents to change belief in a cascade. We show that, over an important range of conditional probabilities, the frequency distribution of the sizes of cascades of belief change (referred to as *avalanches*) in response to a single new data item follows a power law, a key feature of SOC's. Specifically, the distribution of avalanche sizes is dominated by many small avalanches and exponentially fewer large ones. Another key feature of SOC's is that the critical behavior is not dependent on finely tuned parameters, hence we can expect this criticality to occur often, in real-world systems. The power law suggests that large avalanches are relatively infrequent, however when they do occur, if sparked by incorrect data, the result can be the entire team reaching the wrong conclusion despite exposure to primarily correct data. In many domains such as sensor networks in the military, this is an unacceptable outcome even if it does not occur often. Notice that this phenomena was not revealed in previous work, because the more abstract mathematical models were not sufficiently expressive. Finally, our

simulation results show that the inclusion of humans resulted in fewer and smaller avalanches.

1.2 Self-Organized Criticality

Self Organized Criticality is a property of a large number of many body complex systems identified in the pioneering work of Bak et al, characterized by a power law probability distribution in the size of cascades of interaction between system constituents. In such systems these cascades are frequently small and punctuated by very large cascades that happen much less frequently. Self Organized Criticality (SOC) has been used in an attempt to explain the punctuated equilibrium exhibited by many systems [Olami *et al.*1992, Clar *et al.*1994]. The defining characteristic of systems that exhibit SOC Systems is an attractor in the system state space which is independent of parameter values and initial conditions. Such an attractor, typically called a *critical state*, is characterized by a lack of a characteristic scale in the interactions of system constituents. For a system in a critical state long range correlations exist, meaning that perturbations caused by individual system constituents can have system-wide effects. There are many systems which exhibit criticality, but require fine tuning of a global control parameter to enter the critical state. For a system that exhibits SOC, the system will spontaneously enter a critical state due to the intrinsic interactions within the system, independent of system parameters [Bak *et al.*1983].

Systems which exhibit SOC share certain fundamental characteristics. Chief among these are large numbers of constituents interacting on a fixed lattice. Furthermore, for each unit in the system, the number of its neighbors on the lattice which it interacts with is typically a small percentage of the constituents of the system. There are three primary system influences that lead to systems which exhibit SOC. The first is an external drive which attempts to change the state of the individual. Examples include market forces which influence an individual's stock purchasing decisions in economic models or gravity acting on the particles of a sandpile. The second factor is a resistance of the individual to change. In the economic model this would be the individual's cognitive resistance to purchasing a stock, while in the sandpile, friction would play this role. The last factor is a threshold in the local resistance at which the individual relents and changes (toppling grains in the sandpile or the point at which an individual is satisfied that a stock is a worth-while purchase).

These three factors interact to create the conditions necessary for the characteristic scale-free dynamics of SOC. For the power law dynamics of SOC external events can have temporal effects on the system well beyond the instant at which the external influence acted, which in turn results in long range temporal and spatial synchronization. In a system which exhibits SOC the local resistance is naturally balanced such that, most of the time, most individuals are below the threshold to change and building towards it (they are remembering external events). The result is that

most of the time only a few agents will change and propagate information. However, infrequently, many agents will simultaneously reach the threshold and a massive avalanche will occur. For a system in which the local resistance is too large, the system will quickly dissipate local perturbations. We refer to this as a system with a *static* regime. In the *static* regime local perturbations do not lead to system-wide effects. In systems where, local resistance is too low, local perturbations are quickly propagated and massive avalanches are frequent. We refer to this as a system with an *unstable* regime.

To make the SOC concept clear, consider the concrete example of dropping sand onto different parts of a sandpile. If friction is low corresponding to the first regime, dropping sand on the sandpile would always result in many small scale local "avalanches" of toppling sand. Conversely, if friction is too high, grains of sand will not be able to start a flow and avalanches will never occur. However, when friction is balanced between these two extremes, some parts of the sandpile would initially resist toppling due to the addition of sand. Sand would collect on many different parts of the sandpile until a critical threshold was reached locally where the weight of the sand would exceed the friction forces between grains. The friction results in system "memory" which allows for many parts of the pile to reach this critical state simultaneously. As a result, a single grain of sand acting locally could set off a chain reaction between parts of the pile that were simultaneously critical and the entire pile would topple catastrophically. This is why power laws are typical in the statistics of "avalanches" in SOC systems.

1.3 Belief Sharing Model

Members of a team of agents, $A = \{a_1, \dots, a_N\}$, independently determine the correct value of a variable F , where F can take on values from the domain $\{true, false, unknown\}$. Note that the term agent is generic in the sense that it represents any team member including humans, robots, or agents. We are interested in large teams, $|A| \geq 1000$, that are connected to each other via a network, K . K is a $N \times N$ matrix where $K_{i,j} = 1$, if i and j are neighbors and 0 otherwise. The network is assumed to be relatively sparse, with $\sum_j K_{i,j} \approx 4, \forall i$. There are six different network structures we consider in this paper: (a) Scalefree (b) Grid (c) Random (d) Smallworld (e) Hierarchy and (f) HierarchySLO. A comprehensive definition of networks a-d can be found in [Newman2003]. A *hierarchical* network has a tree structure. A *hierarchySLO* has a hierarchical network structure where the sensors are only at the leaves of the hierarchy. Some members of the team, $H \subset A$ are considered to be humans. Certain members of the team, $S \subset A$ with $|S| \ll |A|$, have direct access to a sensor. We assume that sensors return observations randomly on average every second step. A sensor simply returns a value of *true* or *false*. Noise is introduced into sensor readings by allowing a sensor to sometimes return an incorrect value of F . The frequency with which a sensor returns the correct value is modeled as a random variable R_s which is normally distributed with a mean μ_s and a variance σ_s^2 . Agents that are directly connected to sensors incorporate new sensor readings according to the belief update

equation, given by Equation 1.1, this is the version used to calculate the probability that F is true to incorporate a sensor reading that returns *false*:

$$P'(b_{a_i} = true) = \frac{A}{B + C} \quad (1.1)$$

$$A = P(b_{a_i} = true)P(s_{a_i} = false|F = true) \quad (1.2)$$

$$B = (1.0 - P(b_{a_i} = true))_t P(s_{a_i} = false|F = false) \quad (1.3)$$

$$C = P(b_{a_i} = true)P(s_{a_i} = false|F = true) \quad (1.4)$$

Agents use Equation 1.5 to incorporate the beliefs of neighbors:

$$P'(b_{a_i} = true) = \frac{D}{E + G} \quad (1.5)$$

$$D = P(b_{a_i} = true)P(b_{a_j} = false|F = true) \quad (1.6)$$

$$E = (1.0 - P(b_{a_i} = true))P(b_{a_j} = false|F = false) \quad (1.7)$$

$$G = P(b_{a_i} = true)P(b_{a_j} = false|F = true) \quad (1.8)$$

where $P(b_{a_i})$ gives the prior belief of agent a_i in the fact F and $P(s_{a_i}|F)$ gives the probability that the sensor will return a given estimate of the fact (*true* or *false*) given the actual truth value of the fact. Finally, $P(b_{a_j}|F)$, referred to interchangeably as the belief probability or the conditional probability or *CP*, gives a measure of the credibility that an agent a_i assigns to the value of F received from a neighbor a_j given F . Humans are assigned a much larger credibility than other agents. That is, for $a_h \in H$ and $a_k \notin H$, $P(b_{a_h}|F) \gg P(b_{a_k}|F) \forall_{i,h,k}$. Humans also have a higher latency between successive belief calculations.

Each agent decides that F is either *true*, *false* or *unknown* by processing its belief using the following rule. Using T as a threshold probability and ϵ as an uncertainty interval the agent decides the fact is *true* if $P(b_{a_i}) > (T + \epsilon)$, *false* if $P(b_{a_i}) < (T - \epsilon)$, and *unknown* otherwise. Once the decision is made, if the agents decision about F has changed, the agent reports this change to its neighbors. Note that in our model, neither $P(b_{a_i})$, the probability that F is *true* according to agent a_i , or the evidence used to calculate it, is transmitted to neighbors. Communication is assumed to be instantaneous. Future work will consider richer and more realistic communication models.

1.4 System Operation Regimes

This section gives insight into the relationship between the local parameter regimes discussed in Section 1.2 and the parameters of the belief model introduced in Section 1.3.

When $P(b_{a_j}|F)$, the credibility that agents assign to neighbors, is very low, agents will almost never switch in response to input from neighbors so no propagation of belief chains occurs. This is the *static* regime. The *unstable* regime occurs in the belief sharing system for high values of $P(b_{a_j}|F)$. In this regime, agents take the credibility of their neighbors to be so high that an agent will change its belief based on input from a single neighbor and large chains of belief propagation occur frequently.

The SOC regime occurs for values of $P(b_{a_j}|F)$ between the two other ranges. Within the SOC regime, most of the time there are small avalanches of belief propagation. Much more infrequently there are large system wide avalanches. The large avalanches occur because in this regime $P(b_{a_j}|F)$ acts as a local resistance to changing belief when receiving a neighbors belief. As described in Section 1.3 for an agent in the *unknown* state, there are thresholds on $P(b_{a_i})$ above which the agent will change to the *true* state, (and vice versa). In certain circumstances given the data received by agents thus far, many agents will simultaneously reach a state where they are near the threshold. In this case a single incoming piece of data to any of the agents will cause an avalanche of belief changes. Figures 1.1 and 1.2 give an example. In the figure P_0 gives the prior on each agent and S_t a sensor reading received at time step t . Figure 1.1 shows a critical state where many agents are simultaneously near the threshold of $P = 0.8$ as a result of sensor readings $S_1 \dots S_9$. Figure 1.2 shows what happens when one additional sensor reading, S_{10} arrives at time $t = 10$. The numbered arrows shows the ordering of the chain of changes as every message passed causes each agent to change its belief since they all need one additional message to rise above the threshold. The result is a system-wide propagation of belief changes.

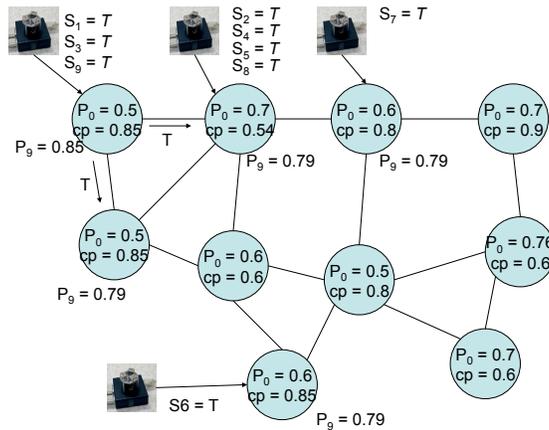


Fig. 1.1. The system in a critical state.

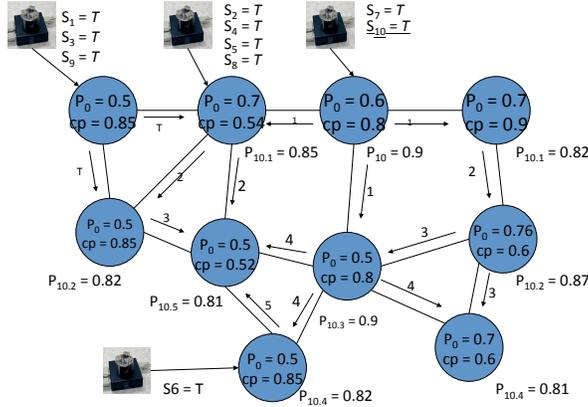


Fig. 1.2. One additional reading causes a massive avalanche of changes.

1.5 Simulation Results

We performed a series of experiments to understand the key properties and predictions of the system. First we conducted an experiment to reveal that the system exhibits SOC. The result of this experiment is shown in Figure 1.3. The x-axis gives the length of a cascade of belief changes and the y-axis gives the frequency with which a chain of that length occurred. The plot is a log-log plot so we expect a power law to be evident as a straight line. The underlying network used to produce the figure was a scale-free network. The 7 lines in the figure correspond to CP values (judgement of a neighbors credibility) of 0.52, 0.54, 0.56, 0.58, 0.6, 0.7 and 0.8. Each point in the graph is an average over 500 runs. While CP values from 0.52 to 0.6 exhibit a power law, values greater than 0.6 do not. This is because if agents have too much faith in their neighbors (large values of CP) then all agents change their beliefs immediately upon receiving input from a neighbor and the system converges quickly. This is the *unstable* regime. While we produced this figure solely for a scalefree network due to space constraints, our experiments show that this trend is true across all networks.

The existence of the SOC regime for a particular range of credibility is important for real world systems because the massive avalanches of belief changes could be propagating incorrect data. Furthermore this occurs rarely and would be difficult to predict. It would be desirable to avoid or mitigate such an eventuality. Understanding the particular parameter responsible and the ranges over which this phenomena occurs allows system designers to reliably do this.

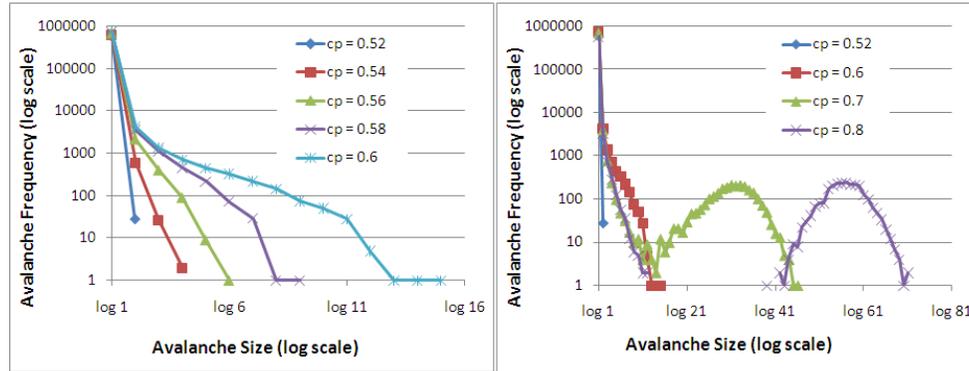


Fig. 1.3. The straight line plots mean that the avalanche distribution follows a power law for the associated parameter range.

Having revealed the SOC regime in the model we wanted to obtain a deeper understanding of the relationship between the parameters of the model and the dynamics of belief propagation in this regime. We concisely capture the dynamics of belief cascades using a metric that we refer to as center of mass. Center of mass is defined as:

$$\frac{\sum_{i=1toN}(AvalancheFrequency * AvalancheSize)}{\sum_{i=1toN} AvalancheFrequency}$$

and can be thought of as the expected length of a chain of belief changes. The next experiment studies the relationship between CP (assessment of a neighbors credibility) and center of mass across the six different network types. Experiments had 1000 agents and 6 different networks. Results are averaged over 500 runs. Figure 1.4 gives the result of these experiments. The x-axis gives CP and the y-axis captures the averaged center of mass. The figure shows a uniform trend across all the networks that when agents have more faith in their neighbors, the center of masses of the avalanches become larger. However, this effect is much less extreme for the networks that are hierarchies. This is because in a hierarchy there are relatively few paths between random pairs of nodes along which belief chains can propagate (all paths go through the root node for example). The effect is particularly prevalent in the networks like the Scalefree and Smallworld networks, because these networks all have a large number of paths between random pairs of nodes.

Next we conducted an experiment to understand the effect of increasing the number of sensors, and thus the amount of information available to the system, on belief dynamics. The x-axis of Figure 1.5, represents the number of agents that have sensors (out of 1000) while the y-axis represents the center of mass. Each point in the graph represents an average of 500 runs over 6 different networks. For this experiment CP is set to 0.56. The graph shows that as the number of sensors increases the center of mass across all the networks types decreases (or increases by only a small amount). A particularly sharp drop is observed for scalefree networks. This is because agents with sensors have more access to information and are thus more certain

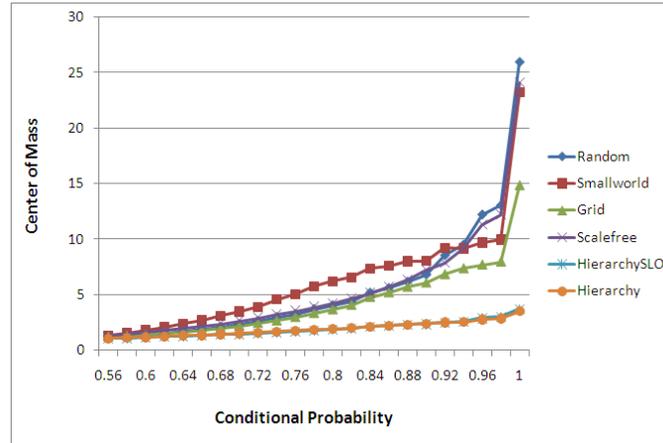


Fig. 1.4. Conditional Probability vs Center of Mass Graph.

in their own beliefs and less likely to change belief in response to information from neighbors. This means that cascades of belief changes are smaller. Conventional wisdom dictates that if information received is primarily accurate then more information is better. We see here however, that increasing the number of individuals that have access to a sensor can discourage the desirable sharing of information. This could be a problem if some sensors in a system are receiving particularly bad data and thus need information from their neighbors.

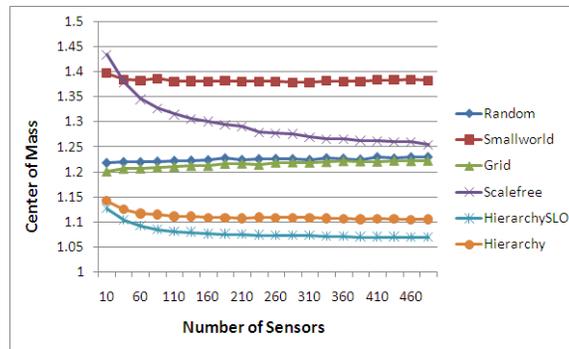


Fig. 1.5. Effect of the increasing the number of sensors on center of mass.

The next experiment looks at the impact of the average network density on the center of mass where average network density = $\frac{\sum a_i K_{i,j}}{|A|}$, which is the average number of links that an agent has in the network. The x-axis of Figure 1.6 shows the average network density increasing from 2 to 8 and the y-axis gives the center of mass. As the

network density increases the center of mass either increases (for scalefree, small-world and random) or remains constant (for grid, hierarchySLO and hierarchy). This is due to the fact that each agent can influence more neighbors and hence lead to the formation of bigger avalanches resulting in a higher or at least equal center of mass.

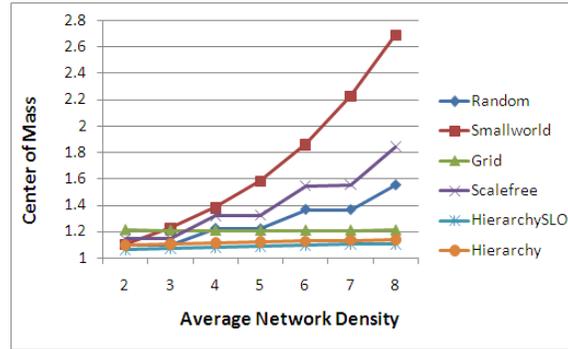


Fig. 1.6. Density vs Center of mass plot.

Next, we conducted an experiment to study the effect that humans have on belief propagation within the system. The model of a human used is described in Section 1.3. Human agents have higher credibility and are slower to change beliefs. The x-axis of Figure 1.7 shows the percentage of humans in the network increasing from 0 to 100% and the y-axis gives the center of mass. Figure 1.7 shows that at first increasing the percentage of humans in a team up to about 20% increases the center of mass of the length of avalanches. The length of avalanches decreases when more than 20% of the agents are human. This effect was noted across all network types. The higher credibility of humans tends to encourage changes in others who receive human input and thus humans can encourage avalanches. Conversely the long change latency of humans has the effect of impeding avalanche formation. When the percentage of humans is below 20% the avalanche enhancing effect of the high human *CP* predominates, while above this value the impedance effect caused by the long latency predominates.

The previous experiments were aimed at understanding the dynamics of belief propagation and the parameters that impact it. However, they did not reveal the nature of the convergence. That is, how many agents actually converge to the correct belief. Our next experiment was conducted to study the impact of network type on the nature of convergence. There are three graphs presented in Figure 1.8 corresponding to random, Small world and Scale free networks respectively. Each graph represents a total of 2500 simulations of the system. The sensor reliability in all the three graphs is fixed to 0.75. A sensor reliability of 0.75 would mean that the sensor gives a true reading with a 0.75 probability. The x-axis gives sizes of clusters of agents that reached a correct conclusion and the y-axis gives the number of simulations

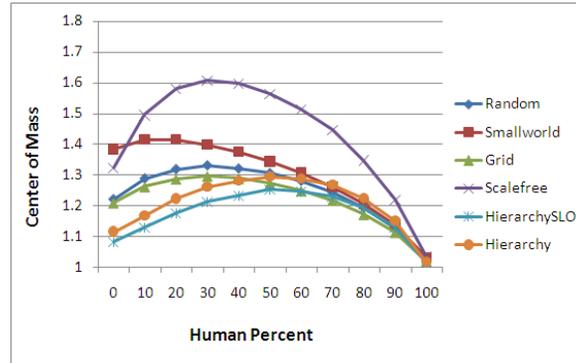


Fig. 1.7. Human percent vs Center of mass plot.

out of the 2500 in which such a cluster occurred. The plot corresponding with the scale free network shows a particularly interesting phenomena. There are two humps in the graph. The larger of the two shows that large numbers of agent converge to the correct conclusion quite often. However, the smaller hump shows that there is a smaller but still significant number of simulations where only very small clusters of agents converge to the correct result. This is a startling result because it says that despite having the same number of sensors with the same reliability the system can come to drastically different conclusions. This is probably due to the nature of the information (correct or incorrect) that started an SOC avalanche which led to convergence. This has important implications in many domains where scale free networks are becoming increasingly prevalent.

The previous experiment gave us insight into the nature of convergence of the system and its relationship to system parameters and network type. However, all previous experiments were conducted using 1000 agents and we wanted to find out if the scale of the system, in terms of number of agents, had any effect on information propagation. In addition, we wanted to understand the sensitivity of the system to changes in system scale, given a particular communication network structure. To this end we conducted an experiment in which the size of the agent team and network type were both varied and the percentage of agents within the system that converged to *true* was recorded. For this experiment 10% of the agents had a sensor up to a maximum of 200 sensors. Sensor reliability was set to 0.75 with $CP = 0.8$, $\epsilon = 0.3$ and $P(b_{a_i}) = 0.5$. Figures 1.9-1.14 give the results of this experiment. For all graphs, the x-axis gives the percentage of agents within the system that converged to *true* and the y-axis gives the number of simulation runs (out of 500) in which the corresponding percentage of agents converged to *true*. Note that there is an implied temporal progression in the graphs. Specifically, if in a particular simulation run the maximum number of agents that converged to *true* was 90%, this implies that there was an intermediate time step within the simulation run at which 50% of the agents had converged to *true*. For each unit on the x-axis the corresponding histogram

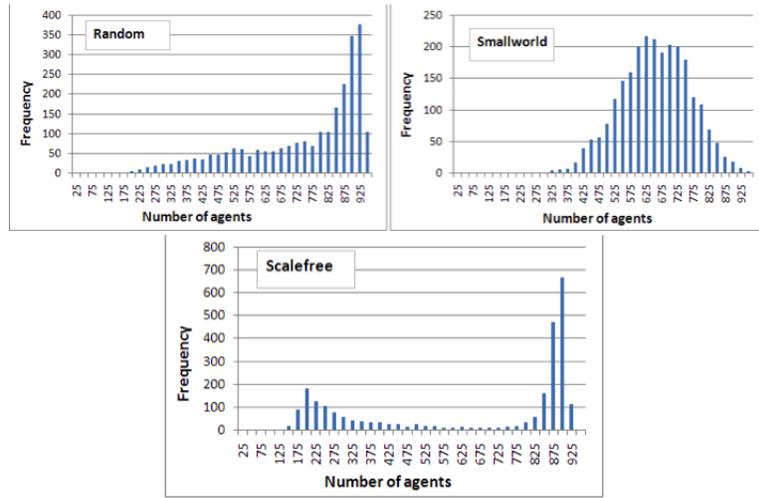


Fig. 1.8. Correct conclusions as a function of network type

bars correspond (from left to right on the plot) to 100, 500,1000,2000, and 5000 agents.

Figure 1.9 gives the results of the scale experiment when the agents communicated using a Grid network (network types are discussed in Section 1.3). This graph shows that scale has a significant impact when a grid is used as a communication network. For the larger systems, with greater than 100 agents, the system infrequently has more than 70% of the agents converging to *true*. This is not surprising because on a grid each agent can only influence a small number of other agents in the system and hence information propagation cascades are easily dampened. What is surprising is that this trend is reversed below 70%. That is below this threshold the larger the system, the more frequently it attains the corresponding percentage of convergence. This cannot be explained by the network structure and must then be a result of the large-scale interaction of the agents.

Figures 1.10 and 1.11 give the results for the two communication networks that are organized as tree structures (Hierarchy and HierarchySLO). Both graphs show that scale has a profound effect on information propagation through a hierarchy. This is because a tree structure isolates large portions of the system from each other. In the extreme case, any information that originates in roughly one half of the nodes in the system must pass through the root node to reach nodes in the other half of the system. This bottleneck acts to dampen the majority of the cascades in the system because the root node is solely responsible for passing this information between the two halves of the system. This means that if the root node makes a mistake, it has a disproportionately large effect. This effect is more pronounced at larger scales

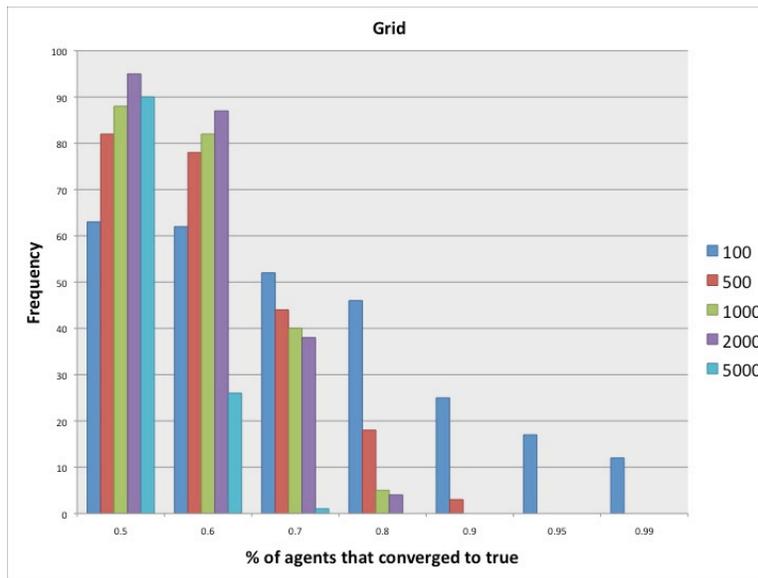


Fig. 1.9. Frequency for which a certain percentage of agents converge to true over different network sizes, The agents use a Grid communication network.

because there is always a single root node in all cases but the size of the isolated sections of the system do scale.

To produce Figure 1.12 a communication network with a Scalefree structure was used. This figure shows, as intuition would dictate, that scale has little effect on convergence for this type of network. This is because the hubs (few agents with high connectivity) act as cascade multipliers, widely disseminating information in the system. However it is surprising that for the Grid structure (Figure 1.9) the system more frequently reaches higher percentages of agents converging to *true*. This can be again explained by the hubs present in the system. Just as in the case of the hierarchy networks, the negative decisions of the hubs have a disproportionately large effect on the system resulting in lower convergence to the correct result.

Figure 1.13 shows the results when a Random network was used. The results are very similar to those for when a Scalefree communication network was employed. This is not surprising because they have similar properties, for example similar average network width. Random connections mean that independent of scale there is a high probability of a short path between any two nodes in the system independent of scale.

Figure 1.14 was produced for a system employing a communication network with the properties of a small-world ring network. The sensitivity of the system employing the small-world ring system to scale seems to share properties with both the Grid-network and the Random network. Like the system employing the Grid network, the percentage of agents converging to *true*, seems to be much higher for 100 agents than

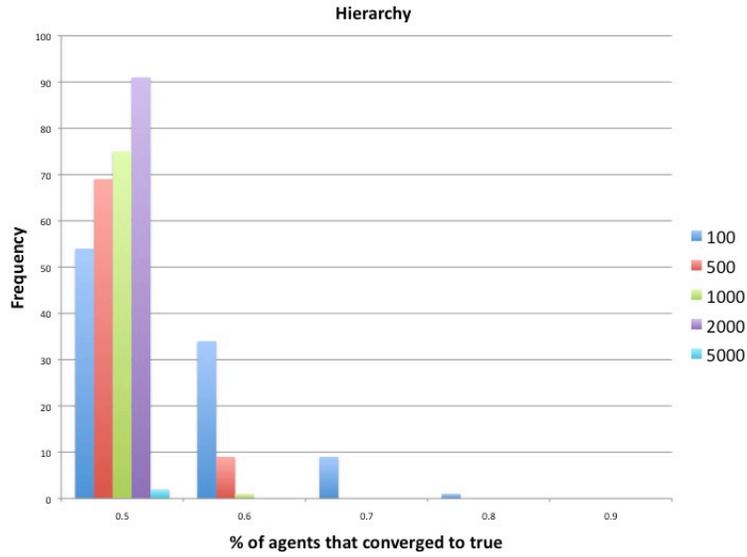


Fig. 1.10. Frequency for which a certain percentage of agents converge to true over different network sizes, The agents use a Hierarchy for communication.

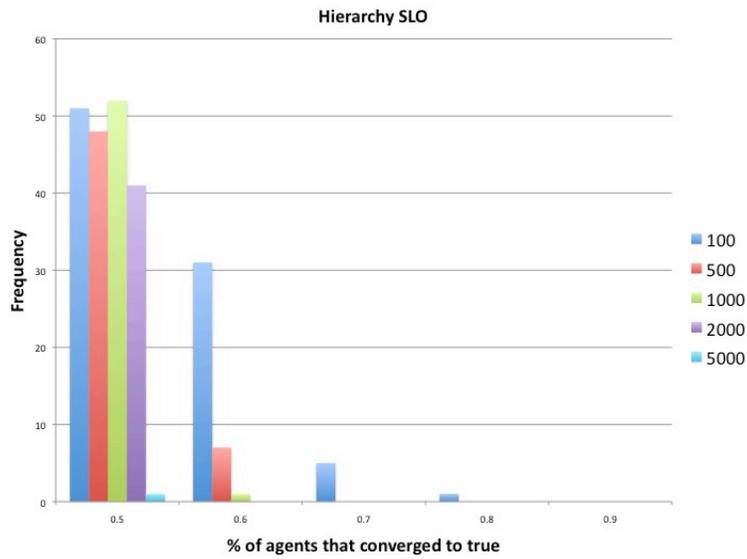


Fig. 1.11. Frequency for which a certain percentage of agents converge to true over different network sizes, The agents use a HierarchySLO communication network.

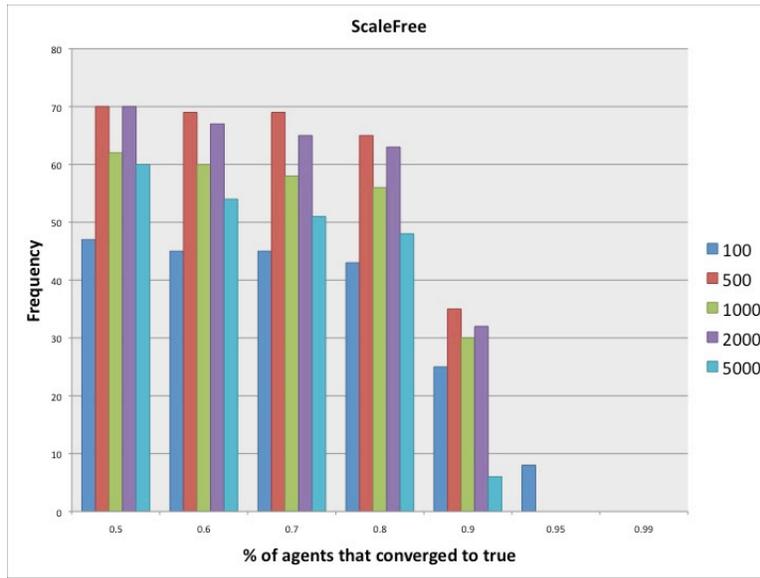


Fig. 1.12. Frequency for which a certain percentage of agents converge to true over different network sizes, The agents use a Scale-free communication network.

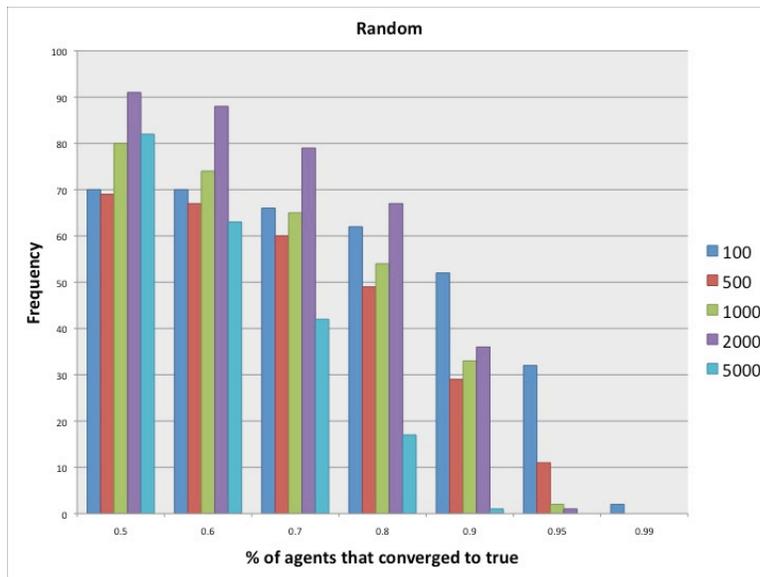


Fig. 1.13. Frequency for which a certain percentage of agents converge to true over different network sizes, The agents use a Random network.

when greater numbers of agents are used. However at other scales the small world ring system behaves in a similar fashion to the Random network. This is reasonable because the small world ring consists of primarily structured connections similar to the Grid network, however, unlike the grid network the small world ring contains of a few random connections.

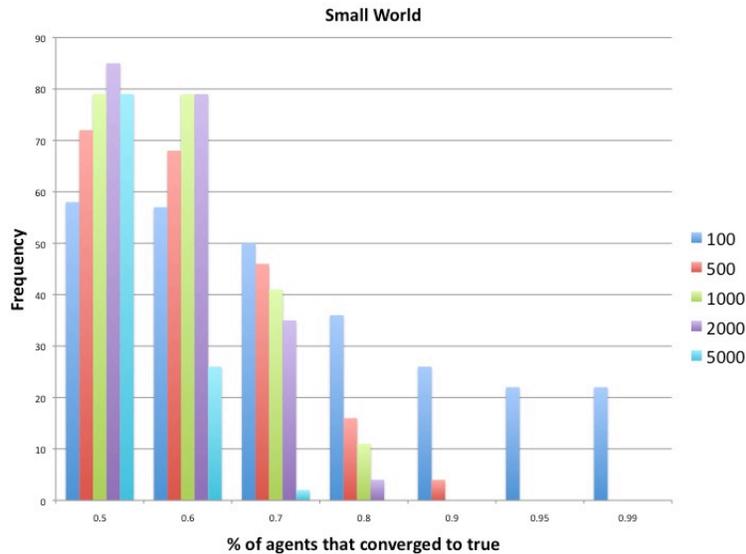


Fig. 1.14. Frequency for which a certain percentage of agents converge to true over different network sizes, The agents use a small-world ring communication network.

1.6 Related Work

Many studies have been conducted to discover systems which exhibit SOC. SOC has provided an explanation for the distribution of earthquake intensity [Olami *et al.*1992]. Drossell showed that forest fire models exhibit SOC dynamics [Clar *et al.*1994]. SOC has been used to explain the avalanche dynamics of a wide range of "pile" models including, sandpiles [Bak *et al.*1983] and rice grains [Malte-Sorensen *et al.*1999]. In addition a number of researchers have attempted to use SOC to describe models of evolution [Zhang1989].

In all of these treatments, simulations of simple models of networked nodes are used to investigate the self-organized criticality of the respective systems. In these treatments the nodes are homogeneous and change their state according to a simple rule that is the same for all nodes. The model described in this paper differs from these others through the use of nodes which are heterogeneous in their decisions about

changing state. The heterogeneity is a result of the distribution of priors across the system. In addition, these distributions of probabilities introduces a dynamic thresholding that is not present in these other models. This means that in the model described in this paper, each node makes a decision about changing state using a unique rule.

The model used in the paper due to Watts [DJ2002] is similar to our own in that it also uses heterogeneous nodes which change state based on a distribution of thresholds to change. There are, however, a number of key differences between our work and that of Watts. The first is that in our model we introduce a simple model of a human. The technical difference introduced by the introduction of human nodes is that some nodes in our system have increased latencies before changing with respect to changes made by their neighbors. Unlike Watts, we explore the effect of these latencies which includes novel system dynamics which fall outside of self-organized criticality. Another key difference is that our model is a three state model (*true, false, unknown*), while Watts uses a two state model (*1,0* equivalent to *true, unknown* in our model). The result of this is an increase in complexity of the decision rules used by nodes which also introduces regimes of system dynamics outside of self-organized criticality. Finally, in the Watts model once an agent changes to the *true* state, it remains in that state for the remainder of the simulation, this is not the case in our model. This has a dramatic effect on the information propagation in the system as well as on the types of emergent effects that are possible. Consequently, in the case of Watts' model only avalanches of *true* states are possible. In [Motter and Lai2002] Motter discusses cascading failures in electrical networks. Motter investigates cascades due to removing nodes as opposed to cascades of decisions as in our work. Unlike the work discussed here, Motter does not examine the relationship between system properties and the distribution and sizes of cascades.

1.7 Conclusions and Future Work

This work demonstrated that a system of belief sharing agents exhibits SOC dynamics for a particular range of agent credibility. SOC dynamics are likely to exist in real-world systems because of the relatively wide parameter range over which this regime occurs. In addition, we discovered that humans can have a dramatic effect on information propagation. Small numbers of humans encourage information propagation in the system while large numbers of humans inhibit information propagation. Finally, we found that for scale free networks, SOC dynamics often propagate incorrect information widely even when much correct information is available.

In the future we plan to investigate the effects of using a more complex model of the belief maintained by an agent, including multivariate beliefs. Of particular interest are multivariate beliefs where variables are conditionally dependent. In real world systems, beliefs are often abstracted as they move up a hierarchy, we plan to include this effect in future belief models and to study the resulting effect on the convergence of the system. We also plan to investigate richer models of the relationship between

agents and the resulting effects on belief propagation, including authority relationships, trust, and reputation. Finally, we plan to investigate the effects of malicious manipulation by an intelligent adversary on the convergence of the system.

References

- [Bak *et al.*1983] P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality: an explanation of $1/f$ noise. *Physical Review Letters*, pages 381–384, 1983.
- [Bellur *et al.*2002] B. Bellur, M. Lewis, and F. Templin. An ad-hoc network for teams of autonomous vehicles. In *Proceedings of the First Annual Symposium on Autonomous Intelligence Networks and Systems*, 2002.
- [Clar *et al.*1994] S. Clar, B. Drossel, and F. Schwabl. Scaling laws and simulation results for the self-organized critical forest-fire model. *Phys Rev E*, 50:10091018, 1994.
- [DJ2002] Watts DJ. A simple model of global cascades on random networks. *Proceedings of the National Academy of Science*, 99:5766–5771, 2002.
- [Malte-Sorensen *et al.*1999] A. Malte-Sorensen, J. Feder, K. Christensen, V. Frette, T. Josang, and P. Meakin. Surface fluctuations and correlations in a pile of rice. In *Physical Review Letters*, volume 83, page 764767, 1999.
- [Motter and Lai2002] Adilson E. Motter and Ying-Cheng Lai. Cascade-based attacks on complex networks. *Physical Review E*, 66:065102, 2002.
- [Newman2003] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167, 2003.
- [Olami *et al.*1992] Z. Olami, H Feder, and K. Christensen. Self-organized criticality in a continuous, nonconservative cellular automaton modeling earthquakes. *Physical Review Letters*, 68:12441247, 1992.
- [Sukthankar *et al.*2008] G. Sukthankar, K. Sycara, J. Giampapa, and Burnett C. A model of human teamwork for agent-assisted search operations. In *NATO Research and Technology Organisation: Human Factors and Medicine Panel Symposium (NATO HFM-142/RTO Symposium)*, B.P. 25, 7 Rue Ancelle, F-92201 Neuilly-sur-Seine Cedex, France, April 2008. NATO RTA / HFM.
- [Zhang1989] Y. Zhang. Scaling theory of self-organized criticality. In *Physical Review Letters*, page 470473, 1989.