Agent-Based Modeling of Cultural Change in Swarm Using Cultural Algorithms

Ziad Kobti
School of Computer Science
University of Windsor
Windsor, ON, Canada N9B-3P4
Email: kobti@uwindsor.ca

Robert G. Reynolds*
Department of Computer Science
Wayne State University
Detroit, MI 48202 USA
Email: reynolds@cs.wayne.edu

Tim Kohler
Department of Anthropology
Washington State University
Pullman, WA 99164-4910
Email: tako@wsu.edu

***Abstract*- The multi-agent Village simulation was initially developed to examine the settlement and farming practices of prehispanic Pueblo Indians of the Central Mesa Verde region of Southwest Colorado [1,2]. The original model of Kohler was used to examine whether drought alone was responsible for the departure of the prehispanic Puebloan people from the Four Corners region after 700 years of occupation. The results suggested that other factors besides precipitation were important. We then proceeded to add economic factors into the simulation, first allowing agents to engage in reciprocal exchanges between kin. This resulted in larger populations, more complex social networks, and more resilient systems. However, the exchange was done randomly and individuals did not remember the transactions. In this paper we explicitly embed the reciprocal exchange process within a Cultural Algorithm, where individual agents can remember individuals that they have cooperated with. Also, in the cultural space the group can learn generalizations about what kind of relative is likely to successfully respond to a request. These generalizations are used to drive changes in requestor behavior. The results of this approach produced an even larger and more complex system exhibiting greater dependence on hub nodes that are sensitive to precipitation.**

I. INTRODUCTION

A. The Village Simulation Project

The multi-agent Village simulation was initially developed to examine the settlement and farming practices of the Pueblo Indians of the Central Mesa Verde region of Southwest Colorado [1,2]. The simulation uses data from archeological site, soil data, and tree-ring data, among others [3,4]. The model is used to explore and enrich our understanding of the region's inhabitants between the years A.D. 900 and 1300. In the original simulation Kohler and his colleagues suggested that something other than precipitation changes had a role in the demographic history of the Mesa Verde Region since the model did not generate the expected depopulation of the region. In follow-up to Kohler's research, initial work by Reynolds and Kobti [5,6,7,8] focused on introducing social networks based upon kinship in order to update the simulation within a cultural framework. In a previous version individuals could exchange resources randomly throughout the kin network. In this version reciprocity is embedded within a Cultural Algorithm [9,10]. Within this framework the individuals can learn from whom to best request resources in the population space. In the belief space the group can learn generalizations about classes of kin that are most likely to be donors within the network. We then examine how the learning improves the performance of the network in this environment.

B. The Cultural Framework

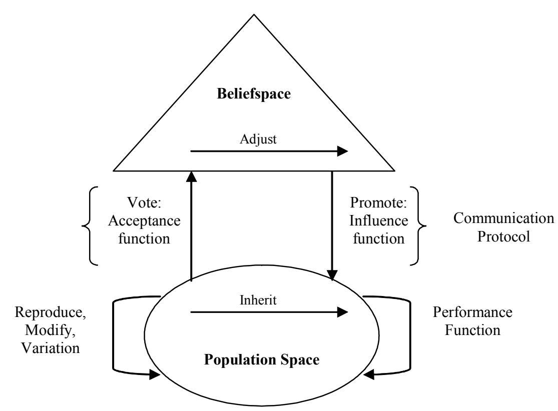

Figure 1. Cultural Algorithm Framework

Cultural Algorithms consist of a social population and a belief space [9] as shown in figure 1. Selected individuals from the population space contribute to cultural knowledge by means of the acceptance function. The cultural knowledge resides in the belief space where it is stored and updated based on individual experiences and their successes or failures. In turn, the cultural knowledge controls the evolution of the population by means of an influence function. A Cultural Algorithm thereby provides a framework in which to accumulate and communicate knowledge so as to allow self-adaptation in both the population and the belief space [11-15].

There are at least five basic categories of cultural knowledge that are important in the belief space of any cultural evolution model: situational, normative, topographic, historical or temporal, and domain knowledge [6]. In our cultural model all of these knowledge sources can be represented. For example, in our current model we assume that agents will have access to knowledge about the distribution of agricultural land (topographic knowledge), the distribution of rainfall as it occurs over time to the extent that it affects agricultural production (history or temporal knowledge), and concerning agricultural planting and harvesting techniques (domain knowledge). These three knowledge sources are fixed at this time. However, situational and normative knowledge can be learned by the group.

C. The Social Network

In previous work [5,6,7,8] we introduced a kinship network into Kohler's model that linked agents with one another. In the model, the basic relationship is based on kinship and the strength of the relationship is impacted by the distance between agents that share a kinship relationship between them. Each agent is a household composed of a husband, a wife and their children. Household members live together in the same location, share their agricultural production, and are affected by the same environmental conditions. Children can grow up, marry, and move out to form their own households. Their connections to their parent households and siblings are maintained in our model. Similarly, the parents maintain ties to their children. When one of the parents in a household dies, the other can form a new household with an available single agent. The initial structure of the social network here supports the notions of parents, siblings, and grandparents on both sides of the family.

The layout of the social network from the perspective of a household is described in table I.

TABLE I: CONNECTED NODES IDENTIFIED BY THE KINSHIP SOCIAL NETWORK.

ParentHHTagA	a link to the parent from the mother's side
ParentHHTagB	a link to the parent from the father's side
ChildHHTag	one link to each child that moves away from this household
RelativeHHTag	one link to each extended family member

The household (agent) rules for marriage and kinship dynamics were described in earlier work [5,9-6]. The social network is defined as the set of all kinship links.

D. Cooperation Framework

Initially, three strategies for reciprocal aid were explored and compared along with the control case in which no goods are exchanged between agents. Table II lists the methods of exchange used and gives a brief description of each. Reciprocal exchange is defined here as exchange of maize between agents related through kinship. Unlike trade between agents, the model of symmetrical reciprocal exchange used does not keep a record of debts owed by particular agents. Modeled after the compassionate and human response of social beings, agents may seek to ask their relatives for food in a time of need, while others donate their surplus to their relatives during prosperous times. In other words the exchange is activated by the requestor, the donor, or both. Each version is potentially reciprocal; the only difference is in terms of who provides the information that triggers the exchange. The current approach to exchange implements a more refined version of Sahlins' (1972) "generalized reciprocity" than that represented previously by Kohler and Yap [16], since the exchanges here are indeed limited to kin, and present as possibilities both asymmetric and symmetric exchanges. Our use of these terms focuses on whether an exchange can be initiated only by the donor or alternatively only by the receiver ("asymmetric") or by either ("symmetric"); symmetric exchange does not imply, here, that the exchanges are balanced in quantity over time.

TABLE II: DESCRIPTION OF THE DIFFERENT COOPERATION METHODS AT THE KINSHIP LEVEL.

Cooperation Method	*Description*
0	No cooperation. No exchange of food.
1	When an agent requires food, it is allowed to select and request food from within its kinship network in order to survive.
2	When an agent has excess food, above a determined threshold amount, it is allowed to select an individual(s) from its kinship network and donate some of its excess.
3	Both methods 1 and 2 are enabled together.

A finite state machine was used to specify the internal states of the cooperating agents. Figure 2 and 3 describe the state model and the transitions between the states.

The states are as follows: Satisfied – An agent is in a "satisfied" state when it has sufficient food in storage to feed the entire household. Philanthropic – An agent becomes a philanthropist when it has a surplus of food in storage, defined in terms of stored maize in excess of a given threshold. For instance, an agent stocking 90% or more of its storage capacity would be able to donate its surplus food. Hungry – A buffer state is implemented at the level just above critical need so that the agent can try to prevent the starvation associated with critical need. When the agent is left with its last food ration then it enters a "hungry" state that triggers precautionary requests for food to avoid starvation. Critical – An agent that has insufficient or no food to eat has no choice but to ask for food or face starvation and imminent death. If the household does not receive its ration to feed the entire family it will die. Death – An agent is marked for immediate removal from the system.

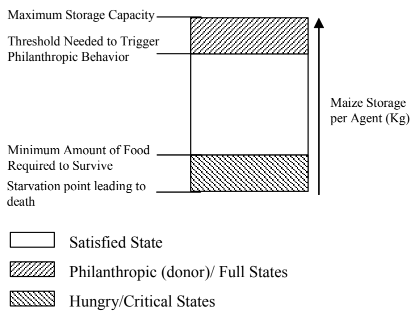

Figure 2: Actual maize amount in storage determines the state of an agent.

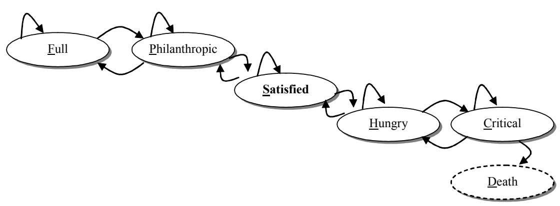

Figure 3: Agent State Transition Diagram. Note that additional states transitions are possible directly between F, P, S, H and C states, but the ones shown are the most frequent.

II. Cooperation Learning Strategies

In previous work [5,6-8] we introduced the cooperation model for reciprocal food exchange across the kinship network. The agent's actions were specified according to a finite state model. In the simulation we can control the agent's capability to cooperate and to what extent. In particular, three cooperative parameters can be specified in addition to no cooperation.

TABLE III: LIST OF THE IMPLEMENTED SELECTION METHODS.

Random	Randomly choose an agent from the kinship network within a given range.
Roulette	Start with random selection of agent (equal initial wheel portion) from the kinship network, then reward or penalize the probability of selecting this agent depending on whether the agents cooperated with or declined the request.

In the initial experiments described in a previous paper [5] the agent used random selection to choose a kin to cooperate (Table III). The constraints were that the cooperating agent has to be directly linked to the initiating agent and within a given radius; we used 4 km. Whether the selected agent actually chooses to cooperate or not was based on that agent's current state and its ability to cooperate with a designated probability. Furthermore, the capability of buffering was added to demonstrate some planning on the part of the individual agent in its struggle to survive. An agent in its hungry state would for instance trigger a request for extra food from its kin as a way to avoid starvation. In a way, we can say the agent preserves its last ration of food for the year in an effort to prevent itself from reaching a critical starvation stage.

A. Kin Selection

The selection of a particular agent for resource exchange was initially random. That, is the agent would always opt to make a random selection within the constraints previously described without keeping track of its past performance. In this paper we now allow agents to maintain plans for the interaction with other agents. Rather than randomly interacting, agents can decide whom they wish to interact with. They now maintain plans concerning whom they prefer to interact with at the individual level, and can produce generalization in the belief space that indicate the type of individuals that are best to exchange with. Thus, every agent possesses and maintains a local strategy for selecting one of its kin for cooperation. These strategies are adapted at the population level using a Genetic Algorithm.

Initially an individual has an equal probability for interacting with each of its kin. Agents are selected for cooperation randomly. However, as time goes on individuals learn to bias the selection process towards individuals who are more likely to interact with them successfully. The selected kin is asked for food based on the requesting agent's need. The selected kin's response is noted by the requesting agent. If the response is successful the requesting agent will maintain the request plan, if unsuccessful it will modify it using the genetic operators. The goal is to allow the agent to identify kin who exhibited a positive response to their request. From an individual's perspective, when it reaches a critical stage and is facing death from starvation, selecting the right donor for food is essential. In the model we can control the number of attempts an agent can make for requests before it gives up and die. In the current experiments the number of attempts was set to two.

B. Individual and Global Strategies

In the previous section we discussed how individuals in the population can learn to modify their plan concerning whom they will interact with so that their interactions are more successful. Learning can also take place at the cultural level in the belief space. In the belief space, generalization about classes of kin who are most likely to respond positively to exchange requests can be produced. What one individual knows can be communicated by others. If an individual discovers that a rich relative is always able to provide food upon request, other related individuals may begin to emulate that behavior. The belief space keeps track of the likelihood that certain classes of kin are more likely to respond to requests than others. This global knowledge is based on the collective experience of the top performers, or exemplars, who report their positive experiences and improve their strategies. When, an individual in the population is about to mutate or change its plans about who to request, the information in the belief space is used to condition the probability of change by making them more likely to select previously successful categories of relatives in order to mutate its existing plan.

C. Learning Methodologies

In the Cultural Algorithm framework, agents learn to adjust their plans at the individual level based upon their experiences and at the cultural level in the belief space. Belief space knowledge is used to condition the changes individuals make to their plans as the result of failures to interact. First we examined the effect of localized learning by individual agents without the sharing of the acquired knowledge with others. Next we introduced situational and normative knowledge in the belief space to track the best performers and the kin types that are most likely to yield successful results in the belief space. Next, the general preferences of the population formed in the belief space are used to influence changes to individual plans in the population space.

Each agent's strategy is comprised of a vector of probabilities for selecting each of its kin as a possible donor (Figure 4). Initially all of its kin have an equal chance to be selected, so a random selection starts the process. The vector contains the likelihood for selecting the mother's household (M), then the Father's (F), then any of its children ($C_1..C_c$), and any of its relatives ($R_1..R_r$).

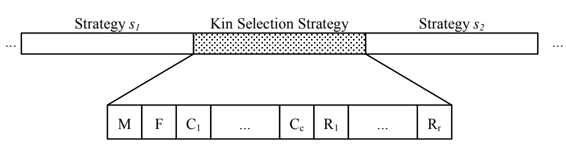

Figure 4. Individual strategy composition for kin selection.

Each individual selects a kin member to interact with using a random process based upon roulette wheel selection. Each possible kin is assigned an area under the wheel that reflects their relative likelihood of selection based upon past performance. The wheel is spun and the selected agent is asked for a donation. If that request is fulfilled successfully by the selected agent then the odds for selecting that kin again will increase. If the request was unfulfilled, then that kin's odds will be penalized and decrease the likelihood of its future selection (figure 5). Over time, each individual agent will learn to adjust its values to reflect its past experiences.

In the next phase, we introduce the cultural belief space as the repository for collecting and using cultural knowledge to guide agent change. In particular two types of knowledge are adjusted dynamically here: Situational and Normative (figure 6). In terms of situational knowledge, exemplars are maintained in the global space to represent individual agents who have been most successful at requesting donations when in need. Everytime an agent completes a request, it updates its local strategy and is evaluated for its maize productivity based on the results of that strategy. In terms of the acceptance of the Cultural Algorithm, the local strategy fitness is then compared to the exemplars currently in the belief space. If the individual's strategy is found to outperform any of the exemplars then it is inserted into the exemplar list. If the maximum number of exemplars in the list is exceeded, the one with the lowest performance score is dropped.

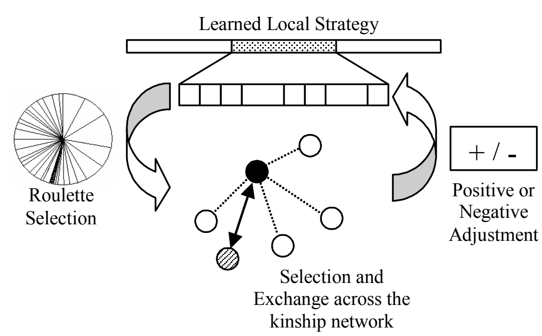

Figure 5. Selection mechanism for an individual agent.

In the belief space Normative knowledge is also used to accumulate information on the frequency with which various kinship types have been successfully selected by exemplars in the population. The ranges reflect the probabilities maintained by the currently selected exemplars. Thus, both normative and situational knowledge can be used to influence the choice of donors by the population. Specifically, if an individual in the population has not been successful in attempting to secure a donor, the knowledge in the belief space can be used to bias its next selection. Using Normative knowledge causes the individual to shift its own probability values closer to the range specified.

Knowledge in the belief space can also be used to influence individual memories as described in figure 7. Each agent, in addition to its local set of probabilities, maintains a local memory that stores the last successful cooperating kin. This list can expand to store additional positive experiences. Currently it is set to one. The method used is to allow the agent to keep the last positive experience in memory so as to be able to use it as the first choice next time it needs to request food. If that agent fails to deliver in subsequent attempts then it is removed from the memory. The individual then selects a new kin to cooperate with from its local strategy conditioned by the culture's Normative knowledge.

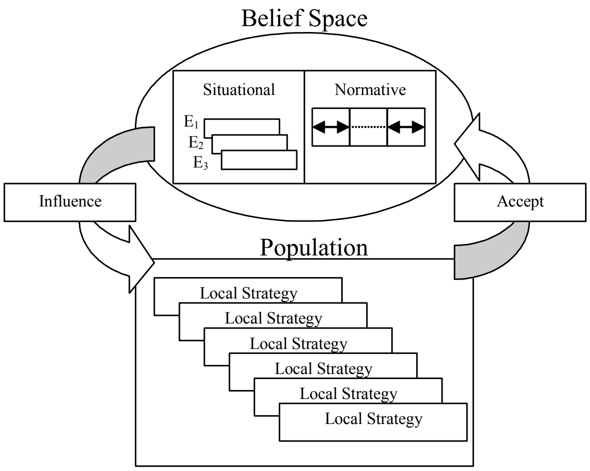

Figure 6. Situational and Normative knowledge in CA.

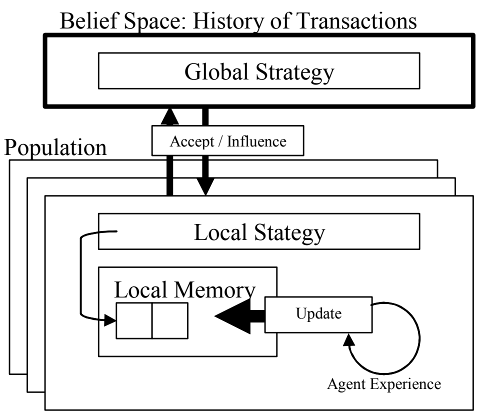

Figure 7. Cooperation learning with memory of last positive cooperating kin.

On a global scale, positive experiences are tracked as exemplars and accumulated in a generalized strategy expressed as normative knowledge. This global strategy can be used to influence the local ones and act as a tracking technique to identify the kins with the ability to provide a positive response when requested for food.

III. RESULTS

In the experiments described here we investigate the impact that learning at the individual and cultural level has on both the system's structural complexity and its resilience. We start with agents making requests to random kin when in need. Next, learning at the individual level is allowed to adjust the probabilities of selecting various individuals. Then knowledge about exemplars and range of probabilities used to select donors is applied to direct changes in individual selection probabilities.

Each of the strategies described above can be used with a different move radius, where the move radius restricts the extent to which agents can relocate and also look for resources. It is viewed as an index of the importance that their current location, or central place, has on their behavior. Move radii in the experiment are set at 10, 20, and 30 pixels (2, 4, and 6 km), respectively. All of the runs shown in figures 8-13 are using a radius of 10.

Figure 8 gives the results of the simulation when there is no cooperation between the agents in the model. It also shows the minimum, maximum and average social links over time for agents (upper graph in the figure) and the social network volume, the product of the out-degree over the total number of the nodes in the kinship network (lower graph in the figure) for this situation. The results are produced between A.D. 900 and 1281 (the horizontal axis in both graphs). Notice that the average number of links for an agent is around 6, and constant over the simulation as is the network volume which is an index of network complexity. Also, there are some nodes with a much higher number of links. We call these hub nodes. These statistics characterize what is often called a small world network, where most nodes have a small number of local links and a few nodes, the hub nodes, have many more general links. These hub nodes serve as the glue that holds the network together.

The scenario distilled in figure 9 allows agents to cooperate, but both potential donors and those in need can request to interact randomly with no memory of past interactions and no learning. However, just the ability to move resources through the network to those in need has an impact on network complexity and variability. For example, even though the average number of links is the same and the maximum size of the number of links is about the same as before, the variability in the maximum size of hub nodes is reduced over that without cooperation. If the variability is in response to environmental perturbations then cooperation ameliorates some of that variation. In addition, the overall network volume is substantially larger, meaning that a larger more complex network can be sustained with cooperation.

In figure 10 we allow learning at the population level, where individuals can adjust their vector of kin selection probabilities based upon experience. Thus, they can remember successful and unsuccessful attempts and use them to adjust the probabilities by random increments accordingly. This results in a slight reduction in variation for maximum hub size, and a slight increase in network volume, although as in previous cases network volume variability increases later in the simulation as a result of periodic drought conditions. Adding learning at the belief level using situational and normative knowledge is shown in figure 11. There the range of probabilities for the selection of each kin type is kept in the belief space. This constrains the adjustment of the individual probabilities in the population space. If for example, an individual is going to increment one of its probabilities but the new value will exceed the upper bound for that kin type in the belief space, the update is constrained to the upper bound.

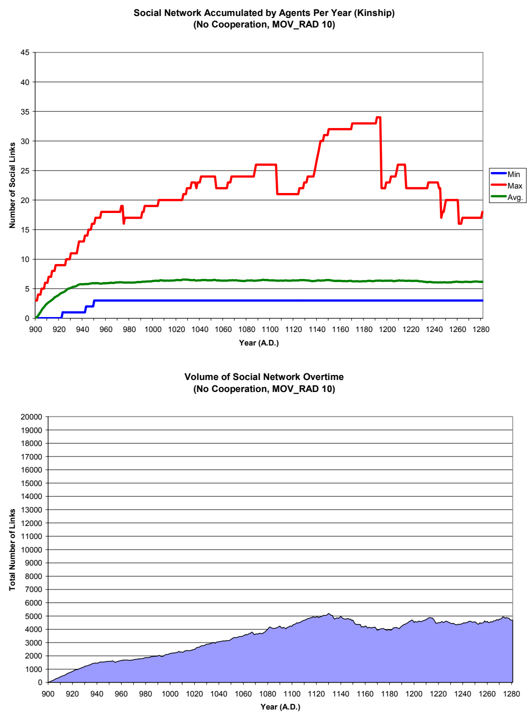

Figure 8. Hub sizes and network volume without cooperation.

The addition of learning in the belief space led to a slight increase in network size overall, and produced an increase in maximum hub size. However, the maximum hub node size exhibited a much steeper decline during the drought phases at the end of the simulation. This trend continues in figure 12 where we show the results of adding the second type of belief space learning to the system. In this case, the belief space knowledge can be used to explicitly select a new exchange partner by replacing the agent's memory when an unsuccessful exchange has taken place. It represents a second way in which mutation can change the plan of the agents in the population. The addition of this type of learning increases the overall network size even more along with the maximum hub size to over 40. Again, this system's network size dips in response to the initial drought phase but recovers to a higher level than with just the one type of belief space learning. As a new learning type is added to the system, the network size increases as well as the maximum hub size. Thus, increased levels of learning produce a more resilient system able to make better recoveries from drought but at the cost of greater dependence on hub node connectivity.

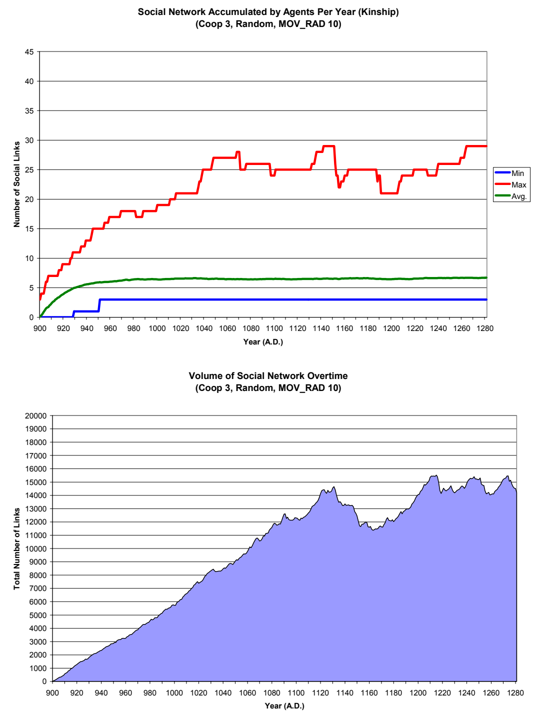

Figure 9. Cooperation with Random selection, no learning.

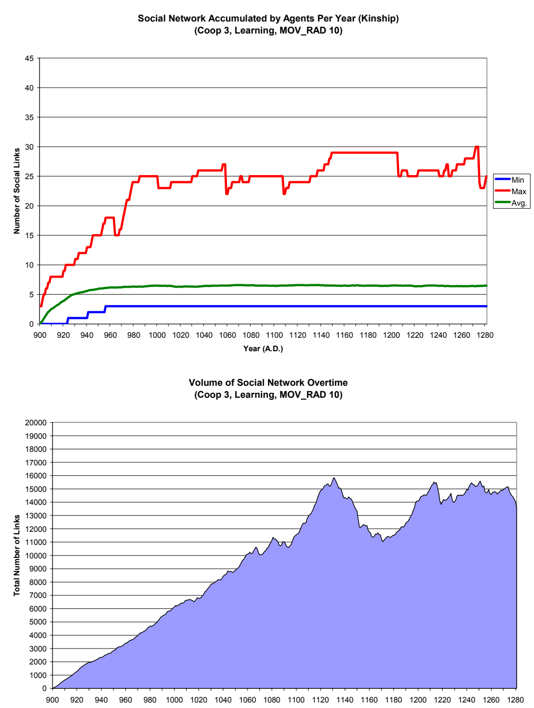

Figure 10. Cooperation with individual roulette learning.

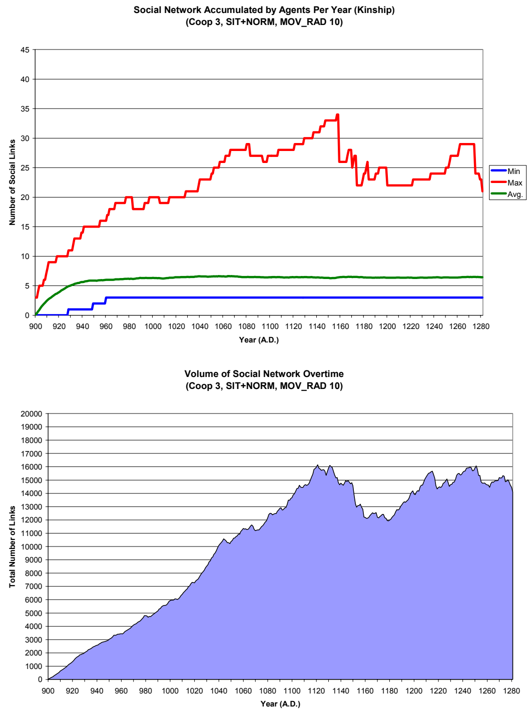

Figure 11. Cooperation with Situational and Normative CA.

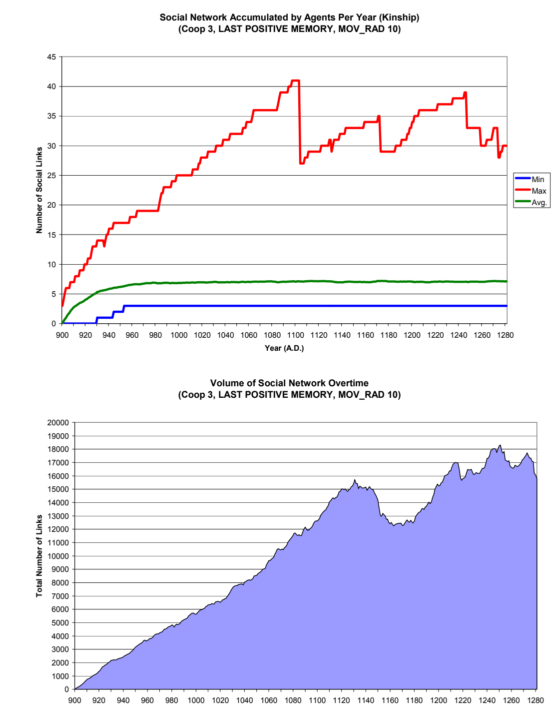

Figure 12. Cooperation with last positive memory.

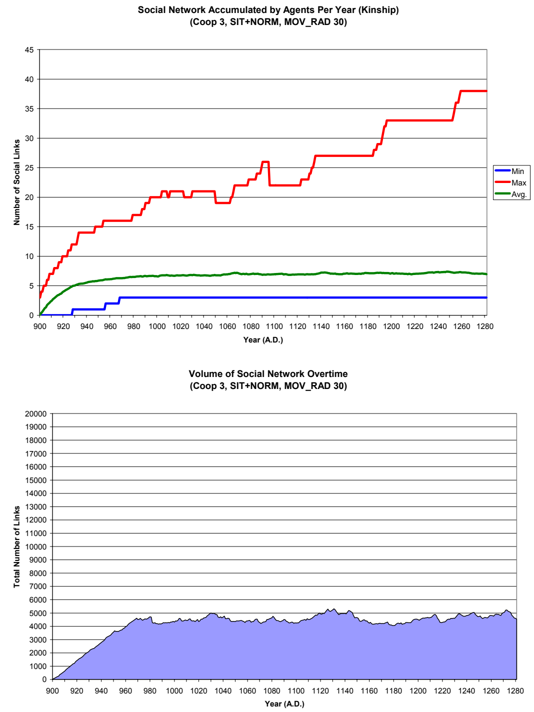

Figure 13. Cooperation with Sit. and Norm. CA, RAD 20.

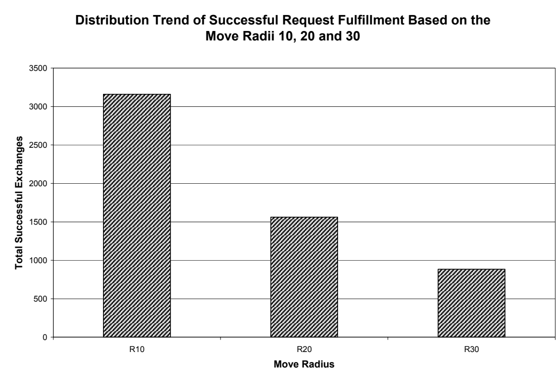

Figure 14. Trend of request fulfillment over various move radii.

Figure 13 demonstrates how important the notion of centrality is to the social organization. If we double the search radius while allowing all three learning activities, the social network is spread out over a larger area and the network size plateaus early at about a quarter of the size when the radius is 10. In figure 14 the number of successful requests are reduced as population density decreases and there are fewer people nearby. This suggests the importance that a social principle of "central place" can have on the social organization. Figure 15 reports on the results of the global memory strategy. Each kin is associated with a frequency accumulated overtime based on the positive requests it fulfilled. The kin with the highest frequency implies that this kin type was the most successful in responding to the needs of others. Notice an almost Poisson distribution with respect to the parent-child relation where the first born has the most links with the parents and it decreases with later children. The second plateau reflects the agent relations with relatives other than its children. They are less intense and of generally equal strength.

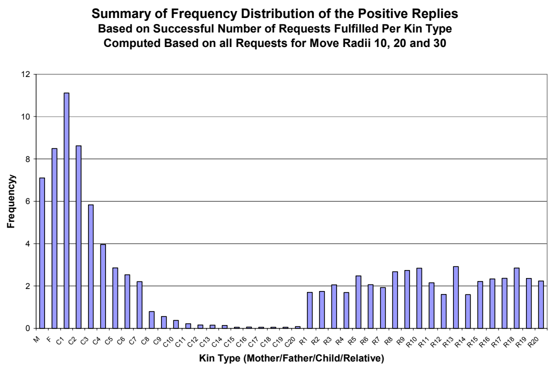

Figure 15. Frequency of kin yielding positive responses.

IV. Conclusions and Future Work

In summary, adding layers of learning to the cultural system allows increase in the system's size and resilience with respect to drought conditions. It bounces back from drought better than when no learning takes place. However, there is a cost that there is more reliance on the hub nodes to "glue" the network together. As long as the hub nodes work in that direction things are okay. But if the hub nodes decide to move as a result of drought conditions how will their movement affect that of the network? In particular, if there is a strong tendency for individuals to aggregate (in other words, if the move radius is small) then if a highly connected individual moves that might cause others in the network to follow. Our results clearly show that individuals could reduce system stress by spreading out over the landscape as an alternative to moving. Thus, a dialectic is presented to the system as to what to do next. Should it reshape to preserve a strong central tendency or spread out instead? In future studies we will look at the impact that resettlement of hub nodes within the landscape will have on the restructuring of the network and the idea of a "central place".

Acknowledgments

This research is funded by NSF grant BCS-0119981.

References

[1] Kohler, T., 2000, "The Final 400 years of pre-Hispanic Agricultural Society in the Mesa Verde Region", Kiva, V. 66:191-264.

[2] Kohler, T., Kresl, J., Van West, C., Carr, E., and Wilshusen, R. H., Be There Then: A Modeling Approach to Settlement Determinants and Spatial Efficiency Among Late Ancestral Pueblo Populations of the Mesa Verde Region, U. S. Southwest. In *Dynamics in Human and Natural Systems: Agent-Based Modeling of Social and Spatial Processes*, edited by T. Kohler and G. Gumerman, pp. 145-178. Oxford, New York.

[3] Van West, Carla. *Modeling Prehistoric Agricultural Productivity in Southwestern Colorado: A GIS Approach*. Department of Anthropology Reports of Investigations 67. Pullman, WA: Washington State University. 1994.

[4] Dean, J. S., and C Van West, Environment-Behavior Relationships in Southwestern Colorado. In *Seeking the Center Place: Archaeology and Ancient Communities in the Mesa Verde Region*, edited by M.D. Varien and R.H. Wilshusen, pp. 81–99. University of Utah Press, Salt Lake City. 2002

[5] Kobti, Z., Reynolds, R., Kohler T., *A Multi-Agent Simulation Using Cultural Algorithms: The Effect of Culture on the Resilience of Social Systems*, IEEE Congress on Evolutionary Computation December 5-12, 2003, Canberra, Australia.

[6] Reynolds, R., Kobti, Z., "The Effect of Environmental Variability on the Resilience of Social Networks: An Example Using the Mesa Verde Pueblo Culture", Society for American Archeology, 68th Annual Meeting, April 9-13, 2003

[7] Reynolds, R., Kobti, Z., "A Multi-Agent Simulation Using Cultural Algorithms: The Effect of Culture on the Resilience of Social Systems", Santa Fe Institute Workshop on Human Environmental Interaction, May 16-18, 2003.

[8] Reynolds, R., Kohler, T, Kobti, Z., *Robustness in Coupled Human/Natural Systems in the Northern Prehispanic Southwest*, Santa Fe Institute/Oxford Studies in the Sciences of Complexity, 2003.

[9] Reynolds, Robert G., *An Adaptive Computer Model of the Evolution of Agriculture for Hunter-gatherers in the Valley of Oaxaca, Mexico*, Doctoral dissertation, University of Michigan, Ann Arbor, 1979.

[10] Reynolds, Robert G., "An Overview of Cultural Algorithms", Advances in Evolutionary Computation, McGraw Hill Press, 1999.

[11] Angeline, Peter A. "Adaptive and Self-Adaptive Evolutionary Computation," in Communication Intelligence, Eds. Marimuthu Palaniswami, et.al., IEEE Press, New York, 1995, pp. 152-163.

[12] Chung, Chan-Jin, 1997, *Knowledge-Based Approaches to Self-Adaptation in Cultural Algorithms*. Doctoral dissertation. Wayne State University.

[13] Conrad, M., 1983. *Adaptability specification as a technique for computing with informal biological models*. BioSystems 13, 303-320.

[14] Flannery, Kent, Marcus, Joyce, and Reynolds, Robert G., *The Flocks of the Wamani*, Academic Press, 1989.

[15] Holland, John H., *Adaptation in Natural and Artificial Systems*, First edition, Ann Arbor: The University of Michigan Press, 1975; First MIT Press edition, 1992.

[16] Kohler, T. and L. Yap, "Modeling Reciprocal Exchange in Southwestern Societies." Paper presented at the 68th Annual Meeting of the Society for American Archaeology, Milwaukee, 2003.