C^2 Adaptation in a Changing Environment

Kathleen M. Carley°
Social and Decision Sciences
and
H.J.Heinz III School of Policy and Management

JuSung Lee
Social and Decision Sciences

Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

Units can and do adapt their C^2 architectures over time and as they engage in operations (combat and peacekeeping). Many of the ways in which units change are not adaptations, i.e., they do not serve to maintain or improve performance. The question is: Which changes represent adaptations? Or in other words: How should the unit adapt to environmental and task changes? Using a computational model of organizational performance, a series of virtual experiments were run to examine which changes in C^2 architecture were adaptations. Using Monte-Carlo techniques the relative performance and change for a set of organizations capable of learning and altering their C^2 structure was measured. The simulated organizations were given a basic set of tasks to perform, then this set of tasks was dynamically switched to another set one or more times. Results indicate that the type of change is more important than the amount of change if organizations are to be adaptive.

1. Introduction

Units often find the need to alter their C^2 architecture in response to changes in the environment. Changes in the environment include such things as certain resources become unavailable or are used up, the specific task being performed changes, the rate at which new problems arise changes, the types of problems faced change, and so on. In response to such situations the unit may alter its C^2 architecture. Indeed, it is often felt that if units do not dynamically respond to a changing environment then unit performance will suffer as the C^2 architecture becomes outmoded. Such changes in the C^2 architecture may be pre-planned as when switching from one well defined architecture to another, or they may be evolutionary as in automatically adapting in response to some local stimuli. Such changes may involve augmenting current staffs with additional personnel, reassigning who is doing what task or has access to what resources, or reassigning who is reporting to whom. However, not all such changes are adaptations; i.e., not all such changes serve to improve performance or even maintain current levels of performance. Some changes, e.g., might interfere with the ability of personnel to learn and to make use of their previous experience, thus degrading the performance of the entire unit.

Within the military context, there may be little time for the CEO to ponder the nuances of the situation and locate the absolutely optimal C^2 solution. Rather, the unit must respond rapidly, make do with what it has, and satisfice. There are typically, in place and often unalterable, constraints (particularly in the short run) that the unit must work around. In other words, the unit is faced with a problem in constrained adaptation. There is little understanding of what types of change represent adaptations, particularly when the unit must operate within constraints.

In this paper, the question of C^2 adaptation in a changing environment is addressed using computational analysis. Before continuing a few caveats need to be mentioned. First, an organization is said to

° This work was supported by Grant No. N00014-97-1-0037 from the Office of Naval Research (ONR), United States Navy.

Appearing in: Proceedings of the 1997 International Symposium on Command and Control Research and Technology. June. Washington, DC

be adaptive if the changes that it makes to its C^2 structure serves to maintain or improve its performance. Second, the organization's C^2 architecture will be characterized as a set of nodes and relations among them. Network based measures will be used to characterize the organization's C^2 architecture. The question of organizational adaptation is addressed by running a series of virtual experiments using a computational model of organizational performance. In this model, ORGAHEAD [Carley, 1996; Carley & Svoboda 1996] organizations have the ability to change aspects of their C^2 architecture and the nature of such changes and their relation to performance is monitored. This model couples a standard model of individual experiential learning with a model of organizations as strategic adaptive agents. At the individual level learning is carried out using a standard experiential learning model based on work in cognitive psychology. At the organizational level adaptation is carried out using a simulated annealing model.

2. Adaptation

Literature on organizational adaptation suggests that organizations do change over time [Stinchcombe, 1965a; DiMaggio and Powell, 1983; Romanelli, 1991]. Part of this change is due to strategic re-organization [Kilman and Covin, 1988] including re-engineering and re-organization. However, not all types of re-organization may be equally valuable. For example, organizational performance may improve as individual members of the organization gain experience [March, 1981]. But, replanning which moves personnel between positions and reassigns them to different tasks may result in loosing the benefits of the experience the members had gained previously. Little is known about organizational change, and even less about how organizations should be designed to promote adaptation. Most theories of organizational design speak to the relative advantage of different designs in different situations [Lawrence and Lorsch, 1967; Hannan and Freeman, 1977]. Such theories, in principle, provide some guidance for organizational change. For example, population ecology can be interpreted as suggesting that if the organization is moving out of a niche environment then the organization should move from a specialist to a more generalist structure [Hannan and Freeman, 1977]. As another example, Staw, Sanderlands and Dutton [1981] have argued that organizations when faced with a decrease in their performance will shift to a more rigid and centralized structure such as is typical in many hierarchical forms. Such suggestions, however, provide little theoretical, let alone practical, guidance as to how to set up the organizations C^2 structure so as to facilitate adaptation. The ultimate goal should be to develop some practical guidelines for adaptation.

Theorizing about organizational adaptation is difficult. The dynamics of change result from simple, but possibly non-linear processes. Consequently, thinking through the implications of adaptation processes is non-trivial. Consider that the following two illustrative processes, retasking and individual learning, may occur simultaneously. When performance drops organizations may enter a downward spiral by choosing to move personnel to work on tasks using resources with which they have little experience, thereby loosing the benefits of the experience these same personnel had in other tasks and with other resources, which in turn may lead to a further reduction in performance, which may lead to further retasking. Alternatively, such retasking may make it possible to perform more tasks simultaneously thereby increasing the speed of organizational response and possibly increasing overall performance regardless of individual expertise or training. Given just these two processes, retasking and learning, what will be the impact of change? Will the retasking be an adaptation? How can issues such as these be addressed?

3. Model

ORGAHEAD is a dual-level information processing model of strategic adaptation in which the unit can change its C^2 architecture in response to various triggers at both the strategic and the operational level. At the operational level units are characterized as being composed of a set of decision making units (DMUs) arranged in some type of command structure. Each DMU may be either a person a subgroup, or a platform. The unit is facing an environment that may change in various ways. The DMUs are adaptive agents, each of whom occupies a particular position and has the capability of learning over time as experience is gained with the task being performed and the resources they are using. Agents are modeled after Bush and Mosteller [1955] as stochastic learning models with additional limits on attention, memory, and information processing capabilities which effectively bound the agent's rationality far beyond those in the original stochastic models. Unit level performance is determined by the actions of the individuals in the organization as they work on tasks. The specific model used is the CORP model of unit performance. At the strategic level, the unit can adapt strategically in response to changes in its performance by altering its design in a number of different ways including personnel movement, retasking, and reassignment of personnel. Unit performance is affected by the ability of the CEO or central unit to anticipate the future and take the appropriate strategic actions to alter the C^2 structure in response to actual or anticipated environmental cues. This strategic adaptation is modeled as a simulated annealing process.

ORGAHEAD has been informed by empirical studies both on individual learning by humans and on adaptation within human organizations. ORGAHEAD is a very versatile program in which the user can specify the initial C^2 structure of the unit (or set of units), whether agents employ SOPs or experience in making decisions, how much training agents receive, how much agents remember, the type of strategic changes allowed, the initial likelihood of the allowable changes, the maximum frequency of change, the rate at which the unit becomes risk averse, the "function" the organization is trying to optimize (e.g., performance or performance subject to minimizing communication), the task environment, and several types of "triggers" for change (such as change in task environment or destruction of resources).

3.1 C^2 Structure and Task Environment

The C^2 architecture is characterized as a series of interlocked networks: authority (who reports to/commands whom), communication (who talks to whom), and resource access (who has access to what resources). Each of these networks can be characterized using various social network measures such as least upper boundedness and span of control. Change in the C^2 architectures can be monitored by examining changes in these measures or in the overall networks.

In the C2 architecture there are 1 to 4 authority levels with 0 to 15 personnel at each level. Each of the agents is boundedly rational and can handle between 0 and 7 resources (including both those needed for communication and for handling the task). As the unit adapts, who reports to whom, who communicates with whom, and who has access to which resources may change. The agent's skill with the resources can vary.

The task is characterized by a nine-bit binary string such that each bit can be thought of as a different, discrete piece of information. In the context of a radar task, the unit responds to a series of yes-no patterns determining whether the pattern represents a friendly or hostile aircraft or vessel. In the context of a situation awareness task, the unit is evaluating a pattern of information to determine the overall nature of the threat. The overall task environment is characterized by its bias and its volatility. The degree of bias is the degree to which the set of tasks faced by the organization all are of the same type or all point to the overall situation as being of the same type. The volatility of the environment is the rapidity with which it switches the set of tasks or types of tasks faced by the units.

3.2 Simulated Annealing

Simulated annealing is a heuristic approach to optimization and is intended to be a computational analog of the physical process of annealing a solid [Kirkpatrick, Gelatt and Vecchi, 1983; Rutenbar 1989]. The goal of the annealing process is to find that state which minimizes costs. The process of annealing involves heating the system to a state that admits many alterations, then, given a schedule of decreasing temperatures, cooling the system slowly so that it reaches thermodynamic equilibrium at each temperature in this schedule, and eventually freezing the system in a good configuration. This process is carried out by having a set of possible moves for altering the existent state to another state, choosing a move, evaluating the proposed state that this move would create, and then moving to that new state if it improves things and possibly even if it does not. Further, the frequency with which such non-improving or risky moves are accepted decreases with time (as the temperature cools). Simulated annealing can be thought of as a computational analog of the process of strategic organizational adaptation through a satisficing process [Carley, forthcoming]. Indeed, based on a detailed empirical study of investment banking Eccles and Crane [1988] argued that the process of strategic change in organizational design gone through by human organizations is an annealing process.

Pictorially one might imagine a three dimensional surface in which all states are arrayed on two dimensions with the height of the surface being the performance or 1/cost of that particular state. The cost function plays a critical role as it determines the shape of the surface. Simulated annealing can be likened to a process of using a mountain climber and a helicopter to locate the highest peak. Imagine randomly dropping the climber on this surface, the climber hikes up-hill, and then the helicopter picks up the climber and puts the climber down in a new location (with decreasing frequency over time). The spot on the surface where the climber is standing is the current state of the system. Like the climber, the system can be in only one state at a time. Clearly, some surfaces, will be easier to traverse and find the highest peak. For example, single peak surfaces are relatively easy to traverse.

Because simulated annealing is a heuristic optimization technique it is not guaranteed to find the optimal solution; nor does it always make the best move. Typically, however, simulated annealing moves the system to a state that is better than where it started. Returning to the pictorial description, if the surface described by the cost function is extremely lumpy then there is no guarantee that helicopter moves will enable the climber to locate the highest point. For combinatorial optimization problems which are NP-complete it may not be possible to locate the exact solution in a reasonable amount of time. Thus, heuristic solutions like simulated annealing are often the only practical answer. We suggest that the organizational design problem is exactly this type of problem. The organizational design problem is the need to locate the C^2 architecture that optimizes organizational performance subject to various constraints. Unit level performance is a function of a large number of factors of which the various elements of design are only one component, but one over which the unit (or at least the CEO) has some, albeit limited, control. Thus, the organizational design problem is, at least, NP-complete. From a purely technical perspective, some type of heuristic based approach appears to be called for.

In ORGAHEAD simulated annealing is used to capture the strategic constraint based adaptation process that the unit goes through. Over time, the unit (more precisely the CEO) attempts to optimize the C^2 architecture given some cost function. The cost function depends on the unit's goal; illustrative goals include maximizing decision accuracy, maximizing kill ratio, minimizing communication. The CEO alters the C^2 architecture strategically; that is, a change is made if it appears to move the organization closer to the goal regardless of whether or not it actually does so [Simon, 1944; March and Simon, 1958]. The organization (more precisely the CEO or central unit) is not omniscient, does not compare all strategies, but simply evaluates a strategy through a kind of "what if" analysis, trying to forecast or anticipate, albeit imperfectly, the future [Allison 1971; Cohen and March 1974; Axelrod 1976]. Since the forecast is known to be imperfect, the CEO may at times gamble on redesigns that might possibly "increase costs" if it is felt that there is some long term advantage. Overtime, the number of high risk moves decreases [Stinchcombe, 1965b] as the unit locks into a certain way of doing business and so gets trapped by its competency [Levitt and March, 1988].

In ORGAHEAD the CEO has a set of possible ways in which the unit's C^2 architecture can be altered. These "ways" include such procedures as change who reports to whom, bring in auxiliary forces, send home forces, and change who has which resources. These "ways" are the move set used by the annealer, and represent the constraints on types of changes the commander can make given legal, economic, availability constraints. The commander proposes a new design (the old design changed by making one of these moves), and then "simulates" the behavior of this possible new design. This is done by giving the hypothetical unit with the proposed C^2 architecture a series of tasks to do. The performance of this hypothetical unit is then compared with performance of the actual unit. In ORGAHEAD the probability of accepting a new design (a strategy from the move set) is defined by the Metropolis criteria and the Boltzman probability criteria. According to this criteria, the commander will always implement the proposed change if the hypothetical structure is expected to be a better performer than the current C^2 architecture. Otherwise, the risky change is accepted with a small probability and that probability reduces over time. The rate of change is set by the temperature cooling schedule. For the virtual experiments that we ran temperature (T) drops each time period as $T(t+1) = a * T(t)$ where a is the rate at which the organization becomes risk averse and t is time.

4. Virtual Experiment

A series of virtual experiments were conducted using ORGAHEAD. These experiments were designed to examine the ability of the unit to adapt given different types of environmental changes. Three kinds of task environments are considered: the no-bias, low-bias, and high-bias conditions. Each bias condition can be thought of as reflecting the likelihood of hostilities. The no-bias condition reflects a state of almost no hostility, or "peace time". Under this condition, the organizational response to the tasks is least critical. In the radar task this means that the criteria for classifying an aircraft as friendly are least stringent. In the low-bias condition, the potential for encountering hostilities is somewhat higher. Thus, in the radar task the likelihood that an observed aircraft is hostile has increased. The high-bias condition can be thought of as "war time". The potential for encountering hostilities is quite high. In the radar task, the likelihood that an aircraft is hostile is very high.

Four kinds of environmental changes are implemented: stable, step, coarse-grained, and fine-grained (see Figure 1). A unit in the stable environment operates in a state of no-bias. That is, the unit faces a series of tasks and the types of tasks remain unchanged overtime and the likelihood of hostilities is relatively low. In the step environment, the set of tasks the unit faces suddenly switches to a different set of tasks. This switch may be from a no-bias to a low-bias condition or from a no-bias to a high-bias condition. The step change occurs half way through the task set, at task 10,000.

The coarse-grained and fine-grained environments oscillate between two states. These are the two basic kinds of environmental changes as posited by organizational ecology [Hannan and Freeman, 1977]. As the name implies, a coarse-grained environment oscillates less frequently than the fine-grained environment. The experiments include oscillations between the no-bias and the low-bias conditions and, also, oscillations between the no-bias and the high-bias conditions. The coarse-grained environment oscillates every 5000 tasks so that the organization response to three changes that are either slight, in the low-bias condition, or drastic, in the high-bias condition over the 20,000 tasks. The fine-grained environment oscillates every 1000 tasks subjecting the organization to nineteen environmental changes over the 20,000 tasks.

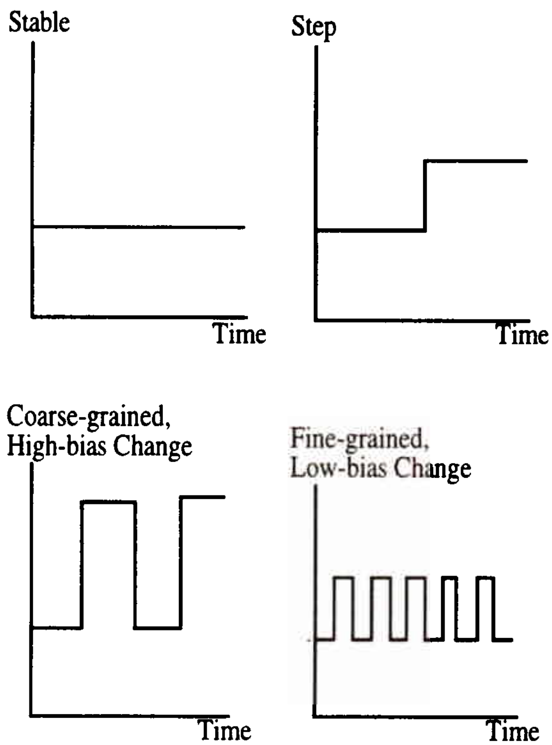

Figure 1. Environmental Changes

In these virtual experiments, a hostile response to a no bias condition occurs when the task pattern contains 5 out of 9 hostile bits. A hostile bit occurs when a bit has certain a numerical value which has been labeled as hostile, in this case 3. A friendly bit is 1. In the low-bias condition, the occurrence of 4 out of 9 bits is sufficient for a classification of hostility. In the high-bias condition, the occurrence of only 3 out of 9 bits in a radar task pattern equates to hostility.

The number of virtual experiments total to 24 (i.e. 3 constraints for adaptation 4 environmental changes 2 biases). Under each experimental condition 1000 units are simulated, using the Monte Carlo technique. Each unit starts with a C^2 architecture (which is characterized by the command structure, who reports to whom, and the control structure, who has access to what resources or does what task). In all of the experiments, for each unit examined, the command and communication structures are equivalent. Each unit's initial C^2 architectures was chosen at random. Each unit is simulated for 20,000 tasks. Changes in the C^2 architecture are allowed to occur every 500 tasks. When the CEO thinks about a proposed change, the "hypothetical unit" is "simulated" for 250 tasks and its performance measured on these 250 tasks. This expected performance is then contrasted with the latest actual performance. Actual unit performance is measured every 500 tasks. When the unit changes its C^2 architecture, it is done so on the basis of the expectation that the new structure will be a better performer, rather than on actual feedback from previous exercises. Thus, some of the changes that are taken may actually, particularly in a changing environment, degrade performance. Each DMU in the unit has a maximum memory capacity of 250 tasks; the DMU behavior reflects what it had learned from solving the last 250 tasks.

5. Measurements

No one measure will suffice The overall C^2 architecture cannot be adequately represented by a single network; i.e., it is more than just the authority relations. Consequently, no one measure of the C^2 architecture will suffice. Thus, we use a suite of measures each of which gets at a different aspect of the C^2 architecture. By examining this suite of measures we will be able to examine various aspects of adaptation.

For each of the 1000 units for each environmental conditions a set of measures were calculated. These measures included information on performance, initial C^2 architecture, final C^2 architecture, and types of changes in the architectures over time. The foremost measurement is performance. Performance is measured as the percentage of last 500 decisions which were accurate. This can be interpreted as the percentage of correct situation assessments or the percentage of aircraft that were correctly classified. Other important measurements include the number of staff augmentations, number of staff sent home or to alternate theaters, number of retaskings and reassignments, span of control, unit level density (fraction of possible authority links that were present), workload for each level in the hierarchy, the number of personnel at each level; and least upper boundedness. Span of control is measured as the average number of subordinates per DMU, or the number of incoming connections. Workload is the number of resources (or equivalently pieces of information) that the DMU must handle. Least upper boundedness [Krackhardt, 1995] is an indicator of the number of levels up that two personnel need to go to resolve a dispute. The lower the level the more decentralized the organization in terms of this type of decision making power.

The average value for each measurement is reported and analyzed as are the average values for the best and worst units (that 5% of the 1000 units with the top/bottom performance). We will refer to those 5% of the units with the highest performance as adaptive units. Similarly, those 5% of the units with the lowest performance will be considered non-adaptive units.

6. Results

The first thing to note is that frequency of change is not the same as adaptation. Indeed, whether we look at stable environments or any other type of environmental conditions, the adaptive units (top performers) make fewer changes in personnel than do the non-adaptive units.

In stable environments, adaptive unit increase in size whereas non-adaptive units decrease in size (Figure 2). While both adaptive and non-adaptive units decrease in density the adaptive units show a greater level of decrease (Figure 3). In other words, adaptation is increasing size and decreasing interconnection by creating a system that decentralizes the organization in terms of resolution of low level conflicts; i.e., least upper boundedness decreases (Figure 4).

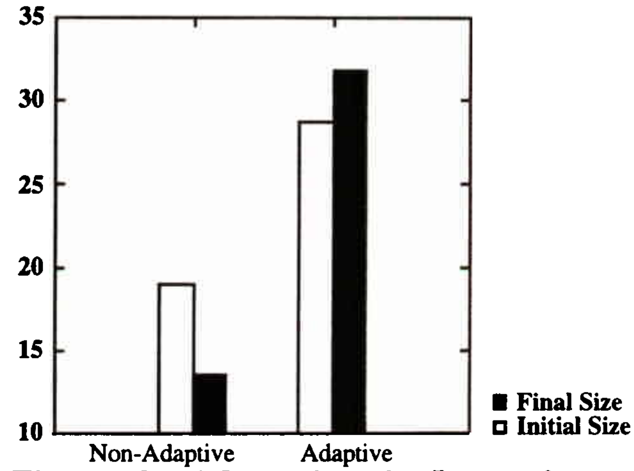

Figure 2: Adaptation is Increasing Size

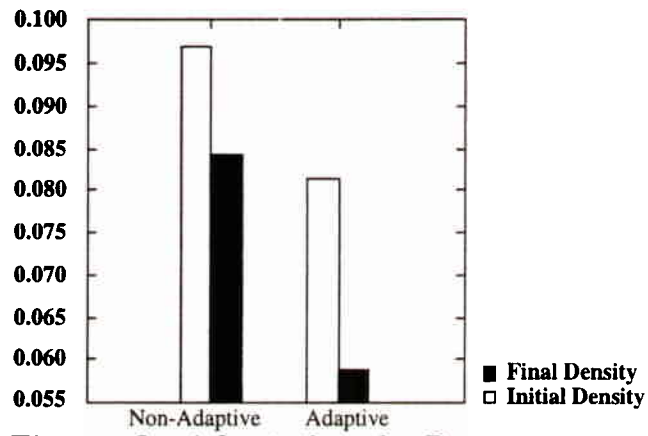

Figure 3: Adaptation is Decreasing Density

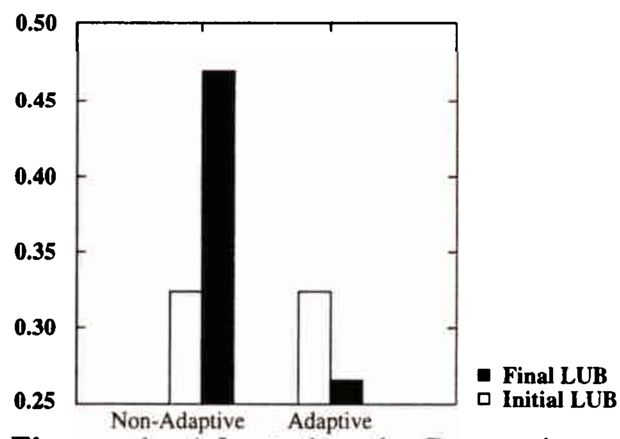

Figure 4: Adaptation is Decreasing Least Upper Boundedness

These same patterns are seen regardless of the environmental conditions. That is, whether the change is coarse-grained or fine-grained or a single step adaptation does not equal frequent change. Interestingly, adaptive units in a fine-grained environment do exhibit more frequent change than do those in a coarse grained environment; though, not as many as in a step environment (see Figure 5). Also in a stable environment the number of changes is lower still. Further, adaptive organizations tend to increase in size and dramatically decrease their density.

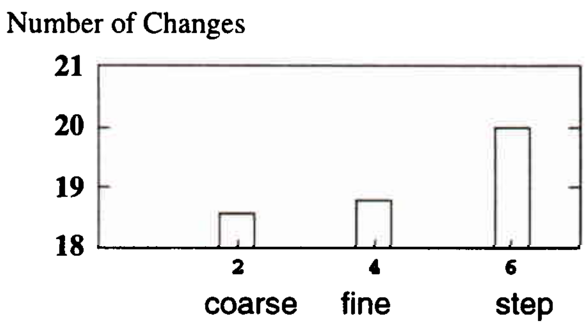

Figure 5: Adaptive Units and Type of Environmental Change

There are interesting interactions between the degree of change and the coarseness of the change. For example, in a coarse-grained environment as bias increases the final size is higher and the final density lower (Figures 6 and 7). Whereas, in a fine-grained environment, bias has no impact on size and the final density is higher under high bias. Finally, while least upper boundedness always goes up with bias, it goes up even more so in a fine-grained environment (Figure 8).

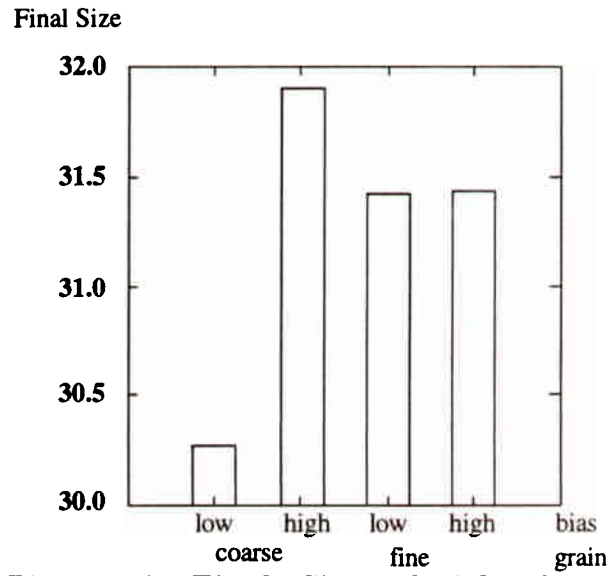

Figure 6: Final Size of Adaptive Units by Environmental Condition

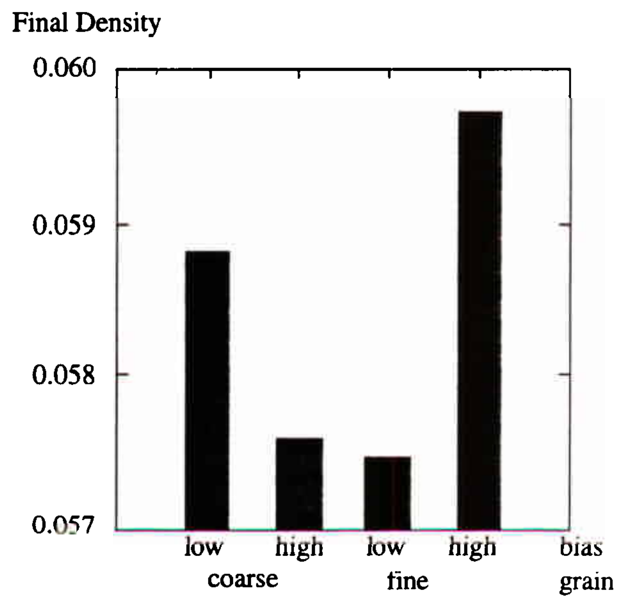

Figure 7: Final Density of Adaptive Units by Environmental Condition

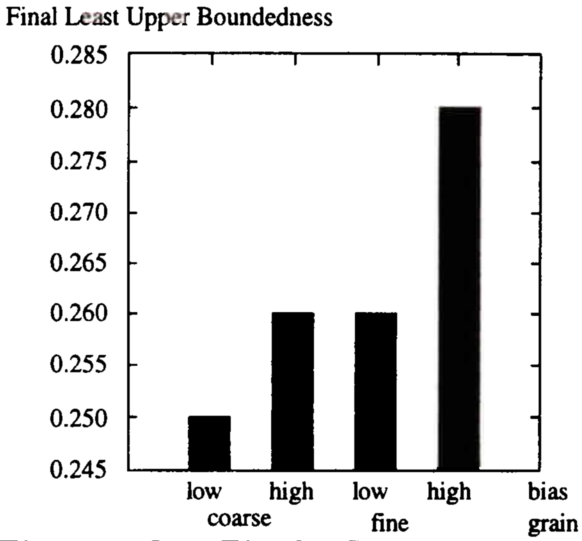

Figure 8: Final Least Upper Boundedness of Adaptive Units by Environmental Condition

7. Conclusions

Using a computational model of organizational performance the ability of units to adapt to changes in the task environment was examined. These results indicate that changing more is not adaptation. Adaptive units actually change less than their non-adaptive counterparts. It is the case that the tougher the environment from a change perspective (more fine grained, or oscillation between very different tasks as in the high bias case) more change is needed; but, the level of change is still low. Further, the pattern of change is a better indicator of adaptation than the amount of change. However, what pattern of change constitutes adaptation depends on the type of task environment. It is well known units that match their C^2 structure to the task structure exhibit better performance. What this research is suggesting is that in a changing environment, units must match the way they change their C^2 structure to the way the environment changes. In other words, the slogans for adaptive units might be: used contingent adaptation, change smarter not more, and when things get tough change more.

References

[Allison, 1971] Graham, Allison, 1971. *Essence of Decision.* Boston, MA: Little Brown.

[Axelrod, 1976] Robert M. Axelrod, *Structure of Decision: The Cognitive Maps of Political Elites.* Princeton, NJ: Princeton University Press, 1976.

[Bush & Mosteller, 1955] Robert R. Bush and Frederick Mosteller. *Stochastic Models for Learning.* New York, NY: Wiley, 1955..

[Carley & Svoboda, 1996] Kathleen M. Carley & David M. Svoboda, "Modeling Organizational Adaptation as a Simulated Annealing Process." *Sociological Methods and Research*, 25(1): 138-168, 1996.

[Carley, 1996] Kathleen M. Carley, "Adaptive Organizations: A Comparison of Strategies for Achieving Optimal Performance" in Proceedings of the 1996 International Symposium on Command and Control Research and Technology. June. Monterray, CA, 1996.

[Carley, forthcoming] Kathleen M. Carley, forthcoming. "A Comparison of Artificial and Human Organizations." *Journal of Economic Behavior and Organization.*, forthcoming.

[Cohen and March, 1974] Michael D. Cohen, and James G. March. *Leadership and Ambiguity: The American College President.* New York: McGraw-Hill, 1974.

[DiMaggio & Powell, 1983] Paul J. DiMaggio and Walter W. Powell. "The iron cage revisited: institutional isomorphism and collective rationality in organizational fields." *American Sociological Review*, 48: 147-160, 1983.

[Eccles and Crane,1988] Robert G. Eccles and Dwight B. Crane. *Doing Deals: Investment Banks at Work.* Boston, MA: Harvard Business School Press, 1988.

[Hannan & Freeman, 1977] Michael T. Hannan and John Freeman. "The Population Ecology of Organizations." *American Journal of Sociology* 82: 929-64, 1977.

[Kilman and Covin, 1988] R.H. Kilmann and T.J. Covin (Ed). *Corporate transformation: revitalizing organizations for a competitive world.* Vol 1. in The Jossey-Bass management series, San Francisco, CA: Jossey-Bass, 1988.

[Kirkpatrick, Gelatt & Vecchi, 1983] S. Kirkpatrick, C.D. Gelatt and M.P. Vecchi. "Optimization by Simulated Annealing." *Science* 220(4598): 671-680, 1983.

[Krackhardt, 1994] David Krackhardt, Graph Theoretical Dimensions of Informal Organizations. In *Computational Organization Theory*, edited by K.M. Carley and M.J. Prietula. Hillsdale, NJ: Lawrence Erlbaum Associates, 1994.

[Lawrence & Lorsch, 1967] Paul R. Lawrence, and Jay W. Lorsch. *Organization and Environment: Managing Differentiation and Integration.* Boston: Graduate School of Business Administration, Harvard University, 1967.

[Levitt & March, 1988] B. Levitt, and J. March, J. "Organizational Learning." *Annual Review of Sociology,* 14: 319-340, 1988.

[March and Simon, 1958] James G. March, and Herbert Simon. *Organizations.* New York: John Wiley & Sons, Inc., 1958.

[March, 1981] James G. March, "Footnotes to Organizational Change." *Administrative Science Quarterly* 26: 563-577, 1981.

[Romanelli, 1991] E. Romanelli, "The Evolution of New Organizational Forms." *Annual Review of Sociology* 17: 79-103, 1991.

[Rutenbar 1989] Rob A. Rutenbar, "Simulated Annealing Algorithms: An Overview." *IEEE Circuits and Devices Magazine* 5: 12-26, 1989.

[Simon, 1944] Simon, Herbert A. "Decision-making and Administrative Organization." Public Administration Review 4: 16-31, 1944.

[Staw, Sanderlands & Dutton, 1981] B.M. Staw, L.E. Sanderlands and J.E. Dutton. "Threat-Rigidity Effects in Organizational Behavior: A Multilevel Analysis." *Administrative Science Quarterly* , 26: 501-524, 1981.

[Stinchcombe, 1965a] A. Stinchcombe, "Organization-creating organizations." *Trans-actions,* 2: 34-35, 1965.

[Stinchcombe, 1965b] A. Stinchcombe, "Social Structure and Organizations." Pp. 153-193 in *Handbook of Organizations* edited by James G. March . Chicago: Rand McNally, 1965.