

LA-UR-6281

*Approved for public release;
distribution is unlimited.*

Title: **DIVERSITY IN DECENTRALIZED SYSTEMS:
ENABLING SELF-ORGANIZING SOLUTIONS**

Author(s): Norman L Johnson Norman@santafe.edu

Submitted to: *DECENTRALIZATION TWO
UCLA NOV 19-20, 1999*

Los Alamos

NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Diversity in Decentralized Systems: Enabling Self-Organizing Solutions

NORMAN L. JOHNSON

Theoretical Division
Los Alamos National Laboratory, MS B216,
Los Alamos, NM 87545

Abstract

Research on how groups can solve problems better than experts has suggested a general understanding of the how diversity plays an important role in a variety of decentralized systems: ecosystems, social groups, large organizations, political systems, free market economies, the stock market, our society - any system where there are individuals or groups make decisions or solve problems without total centralized control or planning. In fact, it appears that diversity is often important in systems thought to be homogeneous and centralized. Diversity is defined to be unique properties of entities, agents, or individuals that are not shared by the larger group, population, structure. Decentralized is defined as a property of a system where the agents have some ability to operate "locally." Both decentralization and diversity are necessary attributes to achieve the self-organizing properties of interest.

Three important observations result from the study that should be helpful to researchers modeling decentralized systems. The first is that dominant processes in these systems can change over time (a developmental view of systems). The second is that self-organization without competition or cooperation is overlooked as an important process, compared to the well-studied actions of competition and cooperation. This type of self-organization can be viewed as a mechanism for achieving higher performance, robustness and conflict resolution without invoking the mechanisms of competition or cooperation for the agents in the system. And finally, measures of robustness of the system also need to be evaluated, in comparison to the often studied measures of performance and efficiency. A consideration of robustness clarifies the role of diversity.

DEFINING CONCEPTS

Because the following text spans many areas of expertise, the following definitions, assumptions and restrictions establish a common perspective for the discussion. We take the approach that the broad variety of systems of interest can be modeled or described as agent-based systems.

An *agent* or *individual* refers to any localized constituent or entity with a decision-making or problem-solving or processing ability. It can be a single individual or a sub-group of agents within a larger system. The decision making or problem solving can be as simple as a deterministic response of a physical subsystem given an initial state and external boundary conditions (because these systems are typically non-linear, deterministic chaos is still possible) or a conscious, premeditative act by a human problem solver. An agent-based model specifies the capabilities of the agent, its environment and the processes by which the agent interacts with the environment and with other agents. A sequence of decisions is a *path* through the problem domain taken by an agent, each step requiring that a previous problem be solved or passed through in order to proceed. For example, a path may be the sequence of decisions by an investor or the foraging of an animal in an ecosystem. An important property of a path is if it has few or many options (*indeterminacy*) in achieving some endpoint. For example, a *robust* (defined below) system often has alternatives paths of equal performance or value; a *fragile* system has few options.

A *group* is a collection of agents that solve a common problem or related problems either knowingly or not, cooperating or not, but which share a common "world" view and expectations within the system. The idea of a common view is discussed later in reference to diversity. *Local* and *global* extent describes the degree of proximity of a property to an agent or group of agents. Local extent is limited to the region of the agent; global extent encompasses the system as a whole. Note that local and global are applied to more than just spatial extent. These concepts apply to any system where the information of the agent is limited to their proximity, including more abstract domains of knowledge space (Johnson, 1998).

A *decentralized* system is where some decisions by the agents are made without centralized control or processing. An important property of agent systems is the degree of *connectivity* or *connectedness* between the agents, a measure global flow of information or influence. If each agent is connected (exchange states or influence) to all other agents, then the system is *highly connected*. Note that connectivity is meaningful only in a decentralized system; if all decisions are made centrally, then there is no significance in the unique information sources of an agent. The degree of connectivity determines how quickly information can flow in the system or equivalently how unique (diverse) the agents can become. In some systems there appears to be an optimal level of connectivity for a system to function (Kauffman, 1993). Note that connectivity can also be a property of environment describing how the landscape is "connected." In most situations, the connectivity of the agent and environment is taken to be the same.

The systems of interest can have the global property of being *self-organizing*, that is, the dynamics are such that the system as a whole exhibits self-regulating or coherent processes that are largely determined by the properties of the subsystems and the governing processes. By definition, a self-organizing system must be decentralized. Global properties that cannot be determined from the properties of the constituents are called *emergent*. For example, the self-regulation of a decentralized capitalistic economy (the invisible hand) is an emergent property of the system (Vriend, 1995).

Note that in the above definitions, the concept of decision making or problem solving is used outside of its normal context of solving a "posed" problem. Problem solving is extended to describe a change of state of a subsystem by agents that may not explicitly pose a "problem," but may just react. The liberty in the definition of problem solving is taken in order to apply a common vocabulary to a variety of systems. This approach is similar to how concepts of cooperation and altruism are applied to both cognitive and non-cognitive systems in biology (Pepper & Smuts, 1999).

Because the systems of interest are often dynamical (non-steady, non-equilibrium), we borrow a few concepts for the study of complex systems. A system is *chaotic* if an infinitesimally small change in the state of some part of the system results in an observable change in the system as a whole. An example in this is a social system is if one vote out of many often changes the outcome or one local battle changes the outcome of a global war. If important properties of the system are insensitive to small changes, then the system is said to be *stable* or *robust* (the opposite taken to be *fragile*). Note that a system can be both robust and chaotic if the chaotic nature does not impact some defined important properties of the system (e.g., the simulations presented in the next section).

Diversity of a group or system is defined to be the degree of unique differences (see (Johnson, 1998) for a mathematical description). Applying this definition, if all the individuals within a group have identical qualities (either experience or capability), then the group has zero diversity, although the qualities of the individuals may encompass all possible variations of the system. If each individual contributes a unique quality not shared by any other, then the diversity of a group is a maximum.

An important consideration associated with diversity is the degree that the agents have a *common world view*, taken to mean that the possible options that an agent have are identical. This does not mean that the preferred options are the same, only that the possible options are the same. The reason that a discussion of diversity must consider world views is because unique qualities are context dependent. For example, the approaches to problem solving of a New Yorker and Australian Bushman are likely mutually exclusive and therefore "unique," but because these approaches operate in very different environments, it is of questionable utility to measure their diversity and ask how it correlates to system performance. This is equivalent to saying that expressions of diversity that can lead to self-organization by the system dynamics require the unique contributions to be

potentially coupled by the system dynamics. Similarly, differences in world view may be the source of conflicts, rather than the diversity itself.

SIMULATIONS OF NON-COMPETITIVE SELF-ORGANIZATION

An idealized self-organizing system which exhibits the features of interest is presented in this section (a summary of a detailed study (Johnson, 1998)). In the understanding presented in the following section on developmental views, the system examined is an idealized *mature* system in which no selection, competition or cooperation is present. It illustrates the mechanisms for diversity creation in mature systems and the importance of diversity to global functionality (performance and robustness). We wish to address the question: what is the simplest demonstration of increased global performance of a group above that of the individual? By most simple, we mean the least number of assumptions, processes or rules.

The idealized system examined is the solution of a sequential problem (Insert in Fig. 1), which has many optimal and non-optimal solutions, solved by agents that have identical capabilities and do not interact. While this maze problem is quite simple from a global perspective, it serves as a representation of more complex processes: the solution of a problem that has many decisions points and many possible solutions (redundancy) and that has difficulty greater than that solvable optimally by one individual. It is argued that a more realistic landscape would not change the underlying processes that are observed in this simple model.

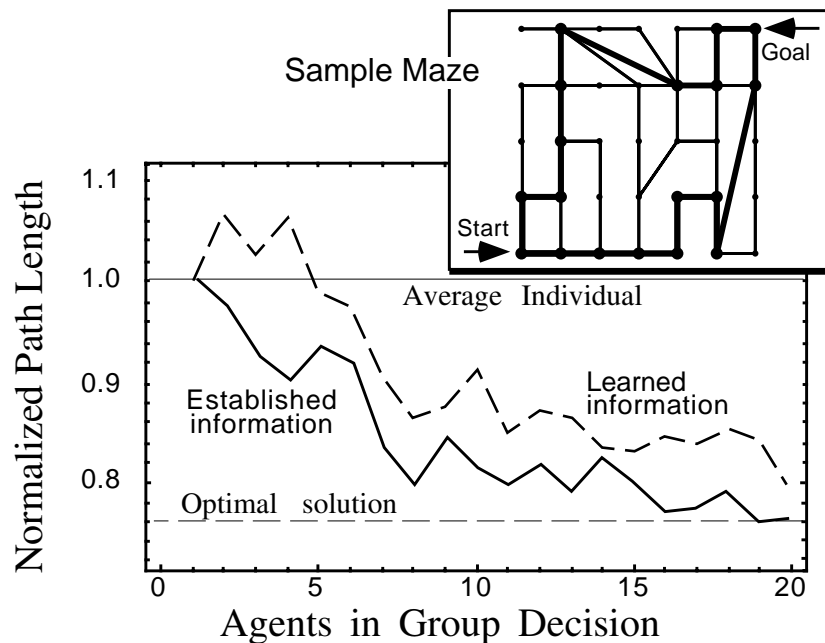


Fig. 1. A sample maze (insert) with two of the 14 minimum paths highlighted and the simulation results (main figure) showing the effect of the group size on the path length, normalized by the average individual path length, about 12.8.

The solution process for a single agent is divided into a *Learning phase* where simple rules of movement are used to explore and learn about the problem domain. Because the agents have no global sense of the problem, they initially explore the problem until the goal is found. The learning process can be thought of as an agent exploring the maze randomly and leaving “breadcrumbs” behind to aid in their search for the goal, thereby avoiding fruitless paths. Then in an *Application phase*, this “learned” information (the bread crumbs) is then used by the agent to solve the problem again, typically with a shorter path as a consequence of eliminating unnecessary loops. Essentially, the agent follows the path with the most breadcrumbs in the Application phase.

Because the initial search is random, a collection of individuals shows a diversity of experience (knowledge of different regions of the maze), diversity of preferences (different preferred paths at any one location in the maze), and diversity of performance (different numbers of steps), even though each agent has identical capability. *This is the source of diversity in the population: by the domain having multiple optimal and non-optimal solutions of equal utility, a diversity of experience, preferences and performance is created.*

In the repeated solution to an unchanging problem domain, we tend to remember only the information that is needed to solve a problem and forget extraneous information associated with unused paths. Here, the equivalent effect is for the agent to remember only “established” information along paths used by individual, thereby “forgetting” unused paths. The process of “forgetting” unused information does not change the performance of an individual agent, because both the learned and established information produces the same path in the Application phase, discounting random choices between paths of equal preference. Therefore, an established individual experience is created from the learned experience by retaining information just used in an individual solution, and forgetting unused information.

Information for a group of individuals is then constructed by a linear combination of the each individual's experiences at each node in the maze. That is, the breadcrumbs from each individual in a selected group are summed for each decision point (node) in the maze. Then the same Application rules as used for the individual are used on this group information to find a group solution. As seen in Figure 1, the group solution always outperforms the average individual for larger groups, and the solution using the established individual information performs better than the learned information. Furthermore, for groups above 20, an optimal solution is found, although nothing in the agent's rules seeks a minimal path length.

Figure 2 shows one mechanism for the reason that the group does better than any agent: individual information is combined to indicate a shorter path for the group. To see how the collective solution is found, pick the path at each node that is traveled by the greatest number of individuals - this corresponds to the maximum preference of the group. Note that the path length of the collective is

better than any individual. The organizational equivalent to this example is how a group of workers may casually share information about their own experience at a certain point in a decision making process (at the water cooler). But because they arrived at a common point from a different path, the relevance of their information may not be apparent and therefore of only of value in a casual exchange and not in direct cooperation. In an organization with many casual exchanges, these seemingly random sources of information can be reinforced by a collective exchange and lead either a group of individuals or a single individual to a better overall solution. This improved collective performance, due minimally to the random social exchange of diverse individuals, also benefits the functioning of the whole. It is also easy to see how this same effect would not happen with a group with low diversity.

The dynamics of the group solution are *chaotic* in detail. For example, the specific path of a group is sensitive to the addition of one individual, even for arbitrarily large groups. Nevertheless, the global solution for the group, any path of minimum number of steps, is stable. This illustrates the desired feature of chaotic dynamics that leads to a responsive and robust system, but not at the expense of the quality of the global solution.




Fig. 2. One mechanism for the better performance of the group. Note that the path length of the group is better than any agent.

To better understand the role and importance of diversity in this simple model, quantitative measures of diversity were examined. The best measure found defines diversity as the degree of unique information in a collection of agents, based on a node-by-node comparison of preferences, as defined in an earlier section. Groups contributing “established” rather than “learned” information exhibit higher diversity, even though less information is available. Moreover, as observed in Fig. 1, the groups based on established information perform significantly better than those based on learned information. Furthermore, this measure of diversity also indicates the degree of insensitivity to noise. In the process of combining information for the group, if valid information of an agent is replaced by random information (breadcrumbs are randomly replaced with some amount), this is a test for the stability of the group solutions. It was found that groups with low diversity were very sensitive to

noise, where groups with high diversity were not: up to 90% of valid information can be replaced before a group path degenerates to a random walk – the worst solution of all methods.

All of the above studies assumed that the agents do not share information while learning or applying information; they are completely independent, except that they solve a problem with a common world view. If the effect of information exchange is included such that the individual can benefit from other agent's experiences while learning the maze, we find that improved individual performance is achieved. Not unexpectedly, groups made up of these shared-learning agents, converge with fewer agents to a minimum path length, much faster than observed in Fig. 1. But, the improved performance is not without a cost. Shared learning results in individuals with similar information and, therefore, the group exhibits low diversity, and, consequently, the stability of the group is degraded, typically severely.

How does the performance of the group depend on the individual performance? Two studies were done, one in which the mazes were made more complex while the individual's capability was held constant, and the other in which the maze was held constant and the individual's capability was varied. From these studies the following conclusions were drawn. 1) A simple maze to a good individual solver is a trivial problem, and no improvement is obtained by a group solution. 2) More difficult global problems require larger groups. 3) An extremely difficult problem to an individual with fixed capability leads to a random individual solution that shows no group advantage. The last conclusion is significant; it suggests that harder and harder problems cannot be solved by larger and larger groups of individuals. Or, equivalently, the individual must have some capability (i.e., not random) which can be amplified in groups.

DEVELOPMENTAL PROCESSES IN DECENTRALIZED SYSTEMS

The main observations of the above non-competitive simulations are threefold. One, they illustrate how diversity can arise in groups of agents *of identical capability* when a system has options of equal likelihood or fitness. Second, they illustrate how random creation of diversity can contribute directly to both global performance and robustness, above that of an individual and in the absence of any selection from the population. Finally, the process of self-organization in the absence of competition and cooperation can duplicate the system-wide advantages of these more traditional properties of decentralized systems. The question addressed in this section is how these non-competitive self-organizing processes can be reconciled with the processes of competition and cooperation. The following is a brief summary of the arguments presented elsewhere for knowledge-based systems (Johnson, 1998) and for ecosystems (Johnson, 1999a).

Paralleling a developmental viewpoint developed for ecosystems (Salthe, 1972; Salthe, 1993), the relative roles of competition, self-organization and cooperation can be distinguished, possibly at the

expense of oversimplification. Three stages are proposed for the development of any decentralized ecosystem, social system and economic system:

- Developing systems (immature systems - such as new ecosystems, emerging economies) are competitive and selective.
- Developed or mature systems are multiply interconnected, symbiotic and robust.
- Senescent (old) systems can be diverse, but are fragile due to rigid interconnections, and can easily return to a developing, competitive system as a result of a change in external conditions or loss of a critical subsystem.

Table 1 presents a variety of properties and estimations of their values for the different systems. Explanations of the choices can be found in ((Johnson, 1999a), (Johnson, 1998)).

Table 1. Comparison of stages of development in Decentralized systems.

Property	Developing	Developed or mature	Senescent
Diversity	Increasing	High	Varied
Interconnectivity Complexity	Low and increasing	High and flexible	Rigid
Chaotic	High at all levels	High locally, low globally	Low
Selection / Competition	High	Low	Low
Self-organization	Low, increasing	High, sustained	?
Source of new diversity	Niche creation	Random selection	None
Individual improvement	High	Little	Little
Group improvement	From individual improvement	From self-organization	?
Decentralization	High	Medium due to formation of global structures	?
Robustness	Varied	High	Low
Scalable	High	High	?
Flexibility	Limited paths	Redundant	Restricted

Some examples of each of these in various systems may clarify the differences. A mature ecosystem is a familiar example of the entries in the middle column in Table 1. A mature ecosystem is composed of diverse species, where each individual living to fulfil their own needs, results in a stable system that benefits all. While competition certainly occurs in mature ecosystems, the global performance (robustness) is due to the non-competitive interactions of a diverse community (Johnson, 1999a). The creation of new diversity is continual, not because of selection as is often argued, but from lack of selection (as in the simulations in the last section). Interconnectivity between species is dynamic (changing), but the system is stable as a whole as long as critical pathways of energy or material flow are not broken. Said another way, the system is locally chaotic (species and individual interactions are unpredictable), but the global system is robust and insensitive to details of the chaotic nature (as in the simulations in the last section). The presence of global processes of self-organization can result in a reduction in apparent decentralization in comparison to the developing systems. Said another way, the greater coupling and self-regulation of the entities in the system duplicates the processes normally associated with centralized control.

Similarly, for social systems (organization or political.), most of the above observations made for ecosystems can be also made (with one major difference, social self-reinforcement, discussed at the end of this section). In particular, the unappreciated aspects of social networks in organizations provide problem solving capability and contingencies that directly result from diverse individuals. A realization of this is that most organizations would fail if employees all thought alike or have little social interactions (in contrast to formal interactions). Mature decentralized economies also have self-regulation (the "invisible hand" of Adam Smith) (Kochugovindan & Vriend, 1998), not just because of the competitive nature of the individuals or even because of direct cooperation, but because of the interdependency of the diverse elements in an economy. By counter-example, it is easy to see that an economy or market that is absent of diversity is quite fragile (only one path for material-energy-monetary flow) and less efficient (serial processing instead of parallel).

Developing systems are most easily observed in economic systems where agents or companies are in extreme competition (e.g., investment capital markets). Social and political systems are similarly in high competition where connections are rapidly changing and outcomes are uncertain. The prediction would be that as these systems evolve more interconnections and dependencies, they would become less competitive as self-organizing processes emerge.

Examples of senescent systems are rare due to their fragility. Very old ecosystems, such as the Australian rainforest that drains into the Great Barrier Reef, are characterized by highly specialized

species, but which only interact with few other species due to highly defensive attributes (e.g., protective poisons and spines). The American automotive industry a decade ago reflected a system that was high evolved but had limited flexibility, few and fixed interdependencies and suffered from over-regulation.

Generally in developing systems or unstable systems in rapid transition, competition between diverse groups occurs, leading to a reduction in diversity through selection - "by survival of the fittest." Through the processes of selection, these systems as a whole improve performance because the individuals that make up the system are selected to perform better. Mature systems, in direct contrast, achieve a higher system-wide performance directly from interactions between diverse populations, without the need for selection (as in the simulations in the last section). If selection does occur, it only retains the status quo; selection of new types is minimal or not at all. In mature systems, then, survival of the fittest becomes "survival of the adequate;" in ecology, it is known as *soft selection* (Wallace, 1970).

Because a common understanding of these self-organizing processes is argued for a variety of systems, a question arises if there are any differences in these processes or additional processes found in some and not in others? One aspect that social systems appear to have, that is absent in other systems, is the ability to have extensive system-wide reinforcement of ideas (paths or dynamics, in general). This effect is a consequence of human mass-communication, whether in ancient times by the social networks or in modern times by the electronic media. This coherence process enables social and economics systems to "self-resonate" to a degree that can reduce diversity and make the overall system less robust. The effect of extreme nationalism on the policies and social diversity of a country is a prime example. This effect can also occur as a positive feedback mechanism, which can cause a larger impact of a change than would otherwise occur. Many stock market crashes are argued to occur by this process.

In the above discussion the assumption is made that the systems of interest in homogeneous in the progress of maturation. Much more likely, natural or human system will have multiple stages of development present simultaneously, particularly as the system undergoes many cycles of maturation and the components of the systems become out of phase. Probably the best example of simultaneous stages of development is in our economic systems, where highly competitive markets can coexist with stable investment markets. Similarly in organizations, competitive dynamics can be simultaneously present with more cooperative structures. The presence of heterogeneous expressions of development of a system may be advantageous to the overall stability of the system.

COMPETITION AND COOPERATION

A strong argument is made in the above text as to the importance and presence of self-organizing processes for improved system performance. Yet, nothing was said about the presence of cooperative interactions in the above discussion, in either ecosystems (as expressed as mutualistic relationships, such as symbiosis) or in human systems (as expressed in cooperative relationships). The viewpoint taken here is that cooperative relationships are structures that develop in systems that duplicate the performance advantages of self-organizing processes, without the associated randomness of these processes. Viewed this way, cooperation becomes the opposite component to selection for system-wide improvement, with self-organization being the neutral alternative between these two extremes.

ROBUSTENESS AS A PERFORMANCE MEASURE

Often modelers focus on better performance or efficiency as appropriate measures of the desirability of a system, particularly in computational systems. The parameters or modeling choices in a system are optimized for these measures. Because a centralized approach often is highest in performance and efficiency and are less complex, models of systems typically result in homogenous agents optimized for the posed problem. It is easy to understand how diversity in this modeling approach is not likely to be considered advantageous. This conclusion is significantly changed if robustness is added to the list of performance measures.

Because robustness is the ability of the system to be stable given some degree of change or noise in the system, static or equilibrium analysis of models do not give insight into robustness. Studies that do consider stability of a model often conclude that diversity of the agents leads to robust dynamics, as observed for a stock market model (Farmer, 1998), decision making systems (Hong & Page, 1997; Hong & Page, 1998), or knowledge systems (Johnson, 1998). One exception to this observation is a study of the stability of repeated n-person games in populations with diverse strategies (Huberman & Glance, 1993). They concluded that the introduction of diversity destabilizes existing coalitions. Why this conclusion differs significantly from the present work is believed to be a consequence of the assumption that the dynamics are determined by essential competitive actions, and not by non-competitive self-organizing processes.

Because robustness is an important property of these systems, it arguably should be one of the performance measures used to evaluate potential models. And because diversity appears to be an essential contributor to robustness, the importance of diversity in decentralized systems should become a common understanding.

CONCLUSION: WHY DECENTRALIZED SYSTEMS?

As a way to summarize the ideas presented above, this section considers the question of what are the pros and cons of decentralized systems. The ideas presented come from a variety of sources ((Arya, Glover & Routledge, 1998; Barnett, Minis & VanSant, 1998; Cloud, ; Foner, 1995), (Johnson, 1998)).

An often cited argument for decentralization is the advantage of parallel processing (Hogg & Huberman, 1993), as opposed to the serial actions of a centralized processor. Said another way, in a proper implementation of a decentralized approach, if the problem domain gets larger, the amount of computation need not increase with size (the question of scalability). The resources of a centralized approach quickly become saturated as a system grows in size, particularly if communication is required to be shared from all the components (here, processing can increase exponentially with the number of subsystems). The critical challenge of implementing a decentralized solution from a previously centralized approach is one of coordinating the independent components, particularly in balancing the "load" among the various subsystems and guaranteeing that the right information is available at the right time and place. Coordination is a common theme of decentralization.

A variety of related advantages also occur in decentralized systems. Localized decision making can make the system more responsive. If there is variation in the environment, then local agents can customize resources to local needs. In social systems where motivation becomes an issue, local control can lead to local satisfaction and higher motivation.

More apropos to the current discussion is how decentralization and diversity complement one another. Decentralization can isolate agents and thereby encourage and incubate the formation of new ideas, function or approaches. Isolation in some form in itself can be a source of diversity (Johnson, 1999b). Particularly in social systems, innovation often requires some degree of isolation in order to strengthen developing ideas to the point that they then can survive exposure to established views or procedures (Mandeles, 1998).

A related advantage of decentralized systems is how conflict resolution can be avoided or be less stressful to the system than in centralized systems. Often in the analysis of centralized systems, particularly from the perspective of game theory, conflict resolution occurs by the direct action of competition or cooperation between agents. In a self-organizing approach, as illustrated in the simulations in an earlier section, the speculation is that the decentralized nature of the system, combined with the dynamics of the system, provides an indirect resolution of conflicting approaches.

Decentralization allows for potential conflict to coexist – a realization that some conflicts arise solely out of being forced to interact, even though the interaction may not be necessary to the overall dynamics. Furthermore, in a self-organizing system, potentially conflicting information can interact through the system dynamics in a way that may lessen the direct conflict. An example is the

pheromone trails of ants: although the individual trails express significant differences in an approach to the problem (finding food), the diverse information is combined in such a manner for the group that there is no conflict between the individuals. In a similar sense the self-organizing maze solutions presented earlier combine diverse information to the benefit of the group, but without the need for invoking competition or cooperation. If this alternative process for conflict resolution is correct, then many of the models that force direct solutions of conflict might be reexamined. This is fully equivalent to the argument made that cooperative-like functionality need not directly arise from the direct expression of cooperation (Hemelrijk, 1997; Hemelrijk, 1998), but can be an emergent property of the dynamics of the system. Hence, potential conflicts may never arise if the agents with the disparate views or approaches are never in contact because of decentralization. And, the dynamics of the system may resolve conflicts without direct confrontation.

Another advantage of a decentralized system is robustness and fault tolerance, as described earlier. These can occur from two sources: from the redundancy or duplication of agent functions or from the contingency potential of a diverse population. The former is an expression that no subsystem is essential to the overall performance. The latter is an expression that breadth of experience enables recovery from potentially destructive paths. Both are observable in the maze simulations described earlier and are typically not associated with a centralized approach.

What are the disadvantages of a decentralized approach? One obvious example is associated with the advantages described above: the reduced efficiency associated with either duplication of resources and the lack of optimal matching of needs and resources.

More subtle disadvantages also occur. A critical challenge in the creation of a decentralized system is determination of the proper forms of communication and coordination that enables the system to function. In evolving systems (nature, economies, societies) these processes naturally evolve to create the necessary self-organizing dynamics. An improperly constructed decentralized system, for example, can result in subcomponents that are under-connected and become isolated. This can lead to miscommunication because the agents may then have different world views.

One major challenge of the creation, or even evaluation, of a decentralized system is the challenge of defining meaningful performance measures. These by their very nature are centralized (global). All the challenges of centralized systems are invoked by this evaluation process.

Because decentralized systems are just now being studied and modeled, there is little understanding of the general processes observed in these systems. Hence, it is difficult to implement a decentralized system without knowledge of these processes. The wise approach to such a challenge is to take advantage of existing self-organizing processes to enable the creation of these systems. For human systems, enabling existing social processes is an example (Johnson, 1999a).

As a final comment, there is a penultimate reason why diverse, self-organizing, decentralized systems are essential to modern approaches to management and organizations, even society. When problems are more difficult than can be solved by an expert (a centralized system) or by a centralized structure (e.g., governments, corporate headquarters, upper-level management), the only alternative is to enable a self-organizing solution (Johnson et al., 1998), either as a resource to a centralized system or a solution method in itself.

ACKNOWLEDGMENTS

The author gratefully acknowledges insightful conversations with Stanley Salthe and many other colleagues that share a common world view. This work is supported by the Department of Energy under contract W-7405-ENG-36

REFERENCES

Arya, A., Glover, J., & Routledge, B. R. (1998, August 1998). Optionality, Decentralization, and Hierarchical Budgeting, http://sulawesi.gsia.cmu.edu/papers/Manager_Option/paper-aug98.pdf.

Barnett, C. C., Minis, H. P., & VanSant, J. (1998,). Democratic Decentralization, <http://www.rti.org/cid/cidpubs.cfm#dem-dec>, <http://www.rti.org/cid/dem-dec.cfm>.

Cloud, D. Profitability Analysis of Strategic Business Structures. <http://www.swcollege.com/acct/morse/ppt/13>.

Farmer, J. D. (1998). Market Force, Ecology, and Evolution, <http://www.santafe.edu/sfi/publications/98wplist.html>.

Foner, L. (1995, Fri Dec 15 09:56:07 1995). Coordination strategies: Degree of decentralization, <http://lcs.www.media.mit.edu/people/foner/Essays/Agent-Coordination/degree-of-decentralization.html>.

Hemelrijk, C. K. (1997). Cooperation without Genes, Games or Cognition. In P. H. a. I. Harvey (Ed.), Fourth European Conference on Artificial Life, (pp. 511-520). Cambridge: MIT Press.

Hemelrijk, C. K. (1998). Spatial centrality of dominants without positional preference. In C. Adami, R. K. Belew, H. Kitano, & C. E. Taylor (Eds.), Artificial Life VI, (pp. 307-315). Cambridge: MIT Press.

Hogg, T., & Huberman, B. A. (1993). Better Than the Best: The Power of Cooperation. In L. N. a. D. Stein (Ed.), SFI Studies in the Sciences of Complexity, (Vol. V, pp. 165-184). Reading, MA: Addison-Wesley.

Hong, L., & Page, S. E. (1997). Problem Solving by Heterogeneous Agents. mimeo <http://ishi.lanl.gov/Documents1.html>.

Hong, L., & Page, S. E. (1998). Diversity and Optimality. mimeo <http://ishi.lanl.gov/Documents1.html>.

Huberman, B. A., & Glance, N. S. (1993). Diversity and Collective Action. In H. H. a. A. Mikhailov (Ed.), Interdisciplinary Approaches to Non-Linear Systems, . Berlin: Springer.

Johnson, N., Rasmussen, S., Joslyn, C., Rocha, L., Smith, S., & Kantor, M. (1998). Symbiotic Intelligence: Self-organizing knowledge on distributed networks driven by human interactions. In C. Adami, R. K. Belew, H. Kitano, & C. E. Taylor (Eds.), Artificial Life VI, . Cambridge, Mass.: MIT Press.

Johnson, N. L. (1998). Collective Problem Solving: Functionality Beyond the Individual, <http://ishi.lanl.gov/Documents1.html>.

Johnson, N. L. (1999a, May 11-14, 1999.). Self-Organizing Knowledge Systems: Enabling Diversity. Paper presented at the 4th Annual Collaboration Conference, Virginia Beach, VA.

Johnson, N. L. (1999b, May 3-5, 1999). Self-Organizing, Collective Problem Solving in Distributed Systems: Functionality without Selection. Paper presented at the Seventh Annual Washington Evolutionary Systems Conference, University of Ghent, Belgium.

Kauffman, S. (1993). The Origins of Order: Self Organization and Selection in Evolution. Oxford: Oxford University Press.

Kochugovindan, S., & Vriend, N. J. (1998). Is the Study of Complex Adaptive Systems Going to Solve the Mystery of Adam Smith's "Invisible Hand"? The Independent Review, 3(1), 53-66.

Mandales, M. D. (1998). The Development of the B-52 and Jet Propulsion: A case study in organizational innovation. Maxwell Air Force Base: Air University Press.

Pepper, J. W., & Smuts, B. B. (1999). The Evolution of Cooperation in an Ecological Context: An agent-based model. In T. A. Kohler & G. J. Gumerman (Eds.), Dynamics of Human and Primate Societies: Agent-Based Modeling of Social and Spatial Processes, . New York: Oxford University Press.

Salthe, S. N. (1972). Evolutionary Biology: Holt, Rinehart and Wilson.

Salthe, S. N. (1993). Development and Evolution: Complexity and change in biology. Cambridge: MIT Press.

Vriend, N. J. (1995). Self-Organization of Markets: An Example of a Computational Approach. Computational Economics, 8(3), 205-231.

Wallace, B. (1970). Genetic Load: Its Biological and Conceptual Aspects: Prentice-Hall.