

LA-UR-99-1819

*Approved for public release;
distribution is unlimited.*

Title: **SELF-ORGANIZING, COLLECTIVE PROBLEM
SOLVING IN DISTRIBUTED SYSTEMS:
FUNCTIONALITY WITHOUT SELECTION**

Author(s): Norman L. Johnson, T-3

Submitted to: *Seventh Annual Washington Evolutionary System Conference,
University of Ghent, Belgium, May 3–5, 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.

Self-Organizing, Collective Problem Solving in Distributed Systems: Functionality without selection (LA-UR-99-1819)

Norman L. Johnson (Ph.D.)

Theoretical Division, MS B216
Los Alamos National Laboratory
Los Alamos, NM 87545
nlj@lanl.gov

Web resource at: <http://ishi.lanl.gov/symintel>

Paper at: <http://ishi.lanl.gov/Documents1.html>

Abstract

To better understand the processes of self-organization and selection in evolutionary systems (ecosystems to social systems to artificial systems), the following discussion focuses on the role in these systems of diversity of groups (composed of agents, individuals, members, sub-groups, etc.) as related to sources of the higher functionality (better performance, efficiency, robustness, adaptability, increased stability, etc.). Diversity was chosen as a clarifying concept, because it appears to be largely ignored or misunderstood, both in its origins and in its contribution to system functionality.

A model system is studied for the collective solution of a sequential problem using self-organizing dynamics. The model system is a large number of non-interacting agents that each solve a maze using simple rules of movement. A collective solution is found by combining individual solutions – linearly. The collective solution exhibits a rich set of properties associated with complex adaptive systems, including emergent properties, redundant subsystems, robust performance, persistent disequilibrium, information condensation, and functionality of the whole greater than the parts.

The increase functionality of the whole is accomplished without any selection of individuals from the population. In fact, it is observed in this model system that any selection results in a decreased performance or stability of the self-organizing, collective solution. This is in contrast to traditional models used for modeling evolutionary systems. The resulting perspective is that selection is a mechanism that increases the functionality of the individual (or group with group selection); self-organizing dynamics of the collective, without need of selection, increases the functionality of the whole above that of the individual. Different expressions of diversity are examined, including individual experience across the domain (breadth) and preferences at one location in the domain (depth). The effect of these on global performance and system stability is examined. This study also provides insight into the question: What individual capabilities are necessary (and therefore have biologically evolved) which enables system-wide self-organization of functionality?

The emergent functionality at a global level makes a strong statement for closure within a bounded system. Unlike other models of evolving systems, it is not necessary to appeal to restrictive assumptions or closure mechanisms to achieve the desired functionality.

Introduction

The following definitions are used:

An agent or individual refers to any localized entity with a decision making or problem solving ability; it can be a single individual or a sub-group of individuals within a larger system.

A collective or group is a collection of agents that solve a problem together, either knowingly or not, cooperating or not, but which share a common view and expectations within the system.

System is taken to be a system of limited extent, closed or bounded in the sense that the agents operate only within the system.

Local and global extent of awareness, system or information distinguishes between a quality of an agent, a variable or a measure. Local extent is limited to the region of the agent. Global extent encompasses the system as a whole.

Diversity within a group is the degree of unique differences within a group (see Johnson, 1998 for a mathematical description). Diversity can either be defined locally or globally. Diversity is a property of a group, not of a single individual. If all the individuals within a group have identical qualities, then the group has zero diversity, even though the qualities may encompass the entire extent of the system. The reason for this specific definition is that it was found to have the best correlation with system performance and stability in an idealized system.

The following main hypotheses are proposed.

1. The primary means for increasing the functionality of the individual is the process of natural selection ("survival of the best fitted"). Diversity of the system is reduced by selection from a diverse population.
2. The primary means for increasing the functionality of the system as a whole is the non-selective interaction of a diverse population of individuals. Diversity is not altered in this process.

The following secondary hypotheses to 2 above is also considered.

- A. Higher diversity leads to higher system performance and robust system dynamics.

Given the correctness of these hypotheses, the resulting perspective is that selection is a mechanism that increases the functionality of the individual; self-organizing dynamics of the system, without need of selection, is a mechanism that increases the functionality of the whole above the ability of the individual. And, the two mechanisms are coupled and mutually evolve.

Background for Natural Selection

The first hypothesis, concerning the role of selection in advancing individual functionality, has been the focus of much evolutionary research and is assumed to be valid. One unresolved issue, which is relevant to the current discussion, is the question of natural selection acting on multiple levels. Group selection has been argued as an explanation of the origin of cooperative or altruistic behavior. The existence of cooperative behavior is based on the argument that if selection can operate between groups, then traits which are a disadvantage to the individual but advantageous to the group (e.g., alarm calling) can be selected and propagated. Prior analytical examinations of this hypothesis required unrealistic assumptions or mechanisms which made the studies unsatisfactory (Ridley, 1996).

A definitive paper on multi-level selection was published recently (Pepper & Smuts, 1999) which presents an agent-based model in an ecological context and essentially addresses the shortcomings of the analytical models. Note that they examine the observed altruistic behavior of alarm callers in predator-prey systems and feeding restraint in a foraging systems. They examine only the amplification of existing traits (mutation is not considered). The authors summarize that "...the model has shown that the groups emerging through the behavior of individual agents in patchy environments are sufficient to drive the evolution of group-beneficial traits, even in the absence of kin selection." They observe that "(b)ecause of its within-group disadvantage, cooperation can only spread though an advantage in founding new groups. Successful groups must be able to export their productivity from the local area, so that their reproductive success is not suppressed by local population regulation." In the absence of "patchiness" or spatial heterogeneity, individuals with the cooperative traits inevitably lose out to their selfish counterparts. Hence, the presence of the spatial extent of the system is an essential aspect of the model. Similarly they argue for the importance of temporal feedback loops that eliminate the need for simplifying assumptions found in the analytic models. "The agent-based approach offers several important advantages, all stemming from the fact that the necessary assumptions concern the traits and behaviors of the individuals rather than the global patterns that arise through their interactions." This is essentially a statement that the global behavior of interest is an emergent property and did not require being "forced" at a global level. The significance of this will be stressed later on the discussion of closure mechanisms.

The variation of within- and between-groups was examined by Pepper and Smuts by looking at different patch sizes and separations of patches. They observed the following.

"Smaller trait groups in turn increased the strength of between-groups selection relative to within-group selection by changing the partitioning of genetic variance. Selection at any level requires that the units being selected vary genetically, and all else being equal,

the strength of the selection increased with the genetic variance among units. In a sub-divided population, all variance among individuals can be partitioned into within- and between group components, and the proportion of the total variance found at each level strongly affects the relative strength of the within- versus between-group selection. The smaller groups are, the more variance is shifted from within to between groups, and thus the stronger the between group component of selection becomes relative to the within group component. Because small isolated patches reduced trait group size, both small patches and large gaps facilitated the evolution of both forms of cooperation."

Within the current context, we conclude from this study that group selection only shifts the observed diversity of the populations to between groups, and reduces the diversity within the groups. Hence, even with the inclusion of cooperation in the selection process, the role of diversity within natural selection as observed in these simulations is essentially as a feed source for selection. In the above example, diversity must be regenerated, either through mutation or migration, otherwise the populations become homogenous. Higher system performance from interdependence is through the interlinked processes, and not through different qualities of interlinked individuals or groups, as will be illustrated next. An example of interlinked processes is that "... *ecological conditions can affect population genetic structure, which in turn can affect the magnitudes of costs and benefits, which in turn can affect reproductive rates and thereby change emigration rates, which in turn feeds back onto population structure.*"

The argument for the origins of diversity in a competitive system

Much has been written to explain the source of the observed bio-, social- and economic-diversity, but little quantitative proof of the mechanisms involved or discussion of the role of diversity have been offered. The basic argument (Rothschild, 1990) is that to minimize utilization of scarce resources, material or energy, an individual will fare better if it can avoid direct competition with other individuals by relocating to an unoccupied niche. By occupying and adapting to new niches, the system as a whole expresses greater diversity. Furthermore, the occupation of new niches also creates additional diversity in individuals that are interdependent with the original relocator. Therefore, the existence of unfilled niches is the driving force in this argument for increased diversity. Hence, the argument requires an unbounded system from the perspective of the individual or group. The combination of the results of the last section and the arguments of this section, therefore, suggest that within a closed system, there is not an argument for increased diversity unless more complex processes are considered.

While the above argument qualitatively explains a reasonable origin of diversity in an extended system, it does

not identify the role that diversity plays on the system as a whole. Certainly one can argue that by increasing the complexity of the interdependencies as a consequence of greater diversity, the system as a whole becomes more robust as additional contingencies are included in the overall dynamics. It is not clear, though, how evolutionary mechanisms would select traits in the individuals for global robustness.

Simulating collective problem solving

In the following text, a quite different system, than the ecosystems considered above, is examined to investigate mechanisms for diversity creation and its importance to systemwide functionality. The following is a brief summary of a much more detailed study (Johnson, 1998).

We wish to address the question: what is the most simple demonstration of increased global performance of a collective 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 Figure 1), which has many optimal and non-optimal solutions, solved by agents that have identical capabilities. While this maze problem is quite simple from a global perspective, it serves as a representation of more complex problems encountered by individuals and organizations: the solution of a problem that has many decisions points and possible solutions and that has difficulty greater than that solvable by one individual. While the system is intended to represent social systems, the relevance to ecological systems is discussed at the end of this section.




Fig. 1. A sample maze (insert) with two of 14 minimum paths highlighted and the simulation results (main figure) showing the effect of the collective 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 a solution is found. The learning process can be thought of as an agent exploring the maze and leaving “bread crumbs” behind to aid in their search for the goal. The Learning rules are a random search with an avoidance of fruitless paths and only use information local to the current position of the agent. 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 bread crumbs. Because the initial search is random, a collection of individuals shows a diversity of experience (different regions of the maze), of preferences (preferred paths at any one location in the maze), and of performance (different numbers of steps), even though they all have identical capability (use the same rules). This is the source of diversity in a population of agents: by the domain having multiple optimal and non-optimal solutions, a diversity of experience, preferences and performance is created, without selection.

In repeated solutions to an unchanging problem, 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. Note that the process of “forgetting” unused information does not change the performance of an agent, because both the learned and established information produces the same path in the Application phase, discounting random choices between paths of equal preference.

Information for a collective of individuals is then constructed by a linear combination of the each individual's experiences at each node in the maze. That is, the bread crumbs from each individual are summed for each path in the maze for a group of agents. Then the same Application rules are used on this collective information to find a collective solution. As seen in Figure 1, the collective solution always outperforms the average individual for larger collectives, and the solution using the established information performs better than the learned information. Furthermore, for collectives above 20, the optimal solution is found, even though nothing in the agent's rules seeks a minimal path length. Figure 2 shows one mechanism for the reason that the collective does better than any agent: individual information is combined to indicate a shorter path for the collective (follow the maximum bread crumbs at each intersection).

To better understand the role and importance of diversity in this simple model, quantitative measures of diversity were examined. The first choice of a measure, the breadth of experience of a collective over all possible paths, was found to be uncorrelated with performance. For example, this measure gives a higher diversity for groups contributing learned experience, even though the experienced groups

perform better. Because of the high variety in the agents' experiences making up the group, this measure of diversity quickly saturates to a maximum value with 10 or more agents.

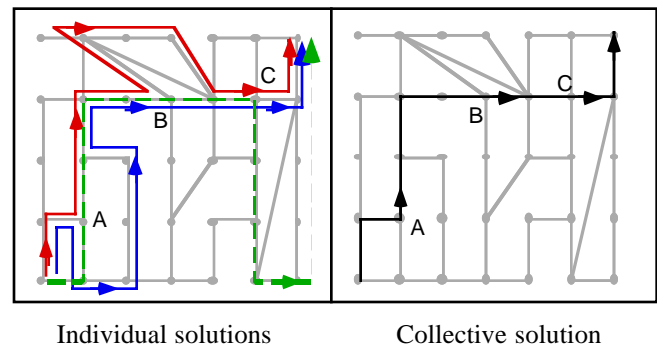


Fig. 2. One mechanism for the better performance of the collective. Note that the path length of the collective is better than any agent.

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 (bread crumbs). For example, if all agents in a group contribute the same information, even if it spans the entire domain, then this measure of diversity is zero. Alternatively, if each agent contributes unique information not shared by others, then this diversity measure is maximum. Interestingly, collectives contributing “established” rather than “learned” information exhibit higher diversity, even though less information is available. And as observed in Fig. 1, the collectives 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 preferences for the collective, if valid information of an agent is replaced by random information (bread crumbs are randomly replaced with some amount), this is a test for the stability of the collective solutions. It was found that collectives with low diversity were very sensitive to noise, where collectives with high diversity were not: up to 90% of valid information can be replaced before a collective 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; they are completely independent. If the effect of information exchange is included in the above simulations, such that the individual while learning the maze can benefit from other agent's experiences, we find that improved individual performance can be achieved, but at the ultimate loss of diversity in the collective. Not unexpectedly, collectives made up of these shared-learning agents, converge with fewer agents to a minimum path length, much faster in Fig. 1. But, the diversity is lost and, consequently, the stability of the collective is severely degraded.

Because the performance of the average individual and the collective were comparable in the above simulations, a question arises: How does the collective advantage depend on the individual performance? This question was addressed by simulations on increasingly larger mazes with up to 2000 individuals in the collective. It was found that the collective solution does not converge to the minimum solution for the most complex mazes when the individual capability is held constant.

A related study was done where global problem difficulty was held constant while the individual capability was varied. From both 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 collective solution. 2) The *rate* of improvement of the collective declines as the maze size is increased; larger numbers of individuals are needed to collectively solve harder problems. 3) More difficult global problems require better problem solvers. 4) An extremely difficult problem to an individual with limited capability leads to a random individual solution that shows no collective advantage. The last conclusion is significant; it suggests that harder and harder problems cannot be solved by larger and larger collectives of individuals with constant capability. It suggests that for a system to solve a problem of high difficulty, a hierarchical approach, domain decomposition or some other mechanism is required.

Before turning to a discussion of the second hypothesis, the applicability of the simulations to non-social systems needs to be addressed. Clearly, the simulations apply to systems where information and knowledge are the landscape for decisions. How can these results be extended to systems such as ecosystems that are dominated by limited resources and that are not obviously sequential in nature, possibly even circular? The limited argument that can be made begins with the observation that ecosystems with fully exploited niches exhibit rich interdependencies that are largely non-competitive and often symbiotic (in the positive sense). These rich interdependencies create many alternative cycles for energy, material and information, in which global system function is not dependent on a single critical path. This is observed in nature as for example in the many sources for transforming sunlight and for decay. The existence of these multiple paths, just as in the above study, leads to a bounded but chaotic system. These systems are chaotic in the same sense as the present study: a small change in initial conditions or by the addition of noise results in a different path through the system. But the system is also bounded in the sense that the global solution is stable, due to the many alternative paths of equal performance.

The effect of information exchange in the Learning phase of the agents is argued to be equivalent to the early formation of a standard functionality in evolving systems. If a predominant approach to adaptation becomes universal, then diversity in this area is low and the system as a whole will become more sensitive to failure in this subsystem,

although the short term performance may be more optimal. An obvious example of this in living systems is the universal adaptation of the DNA/RNA coding system, one of the rare exceptions to biodiversity. If some process in the future should make this adaptation vulnerable, then all life becomes vulnerable.

We conclude that because of dynamics of the above simulations are similar to ecosystems, the results might also be applicable. Admittedly, simulations, which are more directly comparable, are needed to make a stronger argument.

Comparison of the source and role of diversity

The above study illustrates how diversity can arise with agents of *identical capability* within a system which contains multiple options. Just the existence of niches in the problem domain creates diversity. This result occurs without invoking competition to motivate the filling of unoccupied niches, as in the argument for ecological biodiversity based on natural selection. This observation is, in and of itself, not particularly compelling. For an observed source of diversity to be expressed, there must be a parallel need for its existence, and this is discussed next. In general, because of the absence of interaction between agents in the non-selective simulations, one could argue that no exclusion mechanisms exist and therefore, the creation of diversity is solely random.

Given the existence of diversity, the role of diversity in the two systems is quite different. In the argument for natural selection, diversity is an essential requirement for the presence of selection: without diversity, there is nothing to select from. Once the necessary functionality is achieved by selection, the immediate need for diversity is lost. For diversity to remain, the system must be spatially heterogeneous to enable emigration, as in the simulation for group selection, or experiencing changing environment conditions, in which case diversity is needed for future performance of the system. Diversity does not directly contribute to current system performance; only past diversity contributes to the current performance. By direct contrast, the non-competitive system has reduced performance and increased sensitivity in the absence of diversity, independent of future need. With few exceptions in a non-selecting dynamical system, a direct correspondence exists between system performance and the current state of diversity.

In both extremes of selection and non-selection, interdependency plays an important, but different, role. In the study by Pepper and Smuts [1999] the unhindered interdependencies of the *processes* enhanced the system functionality beyond the analytical models. In the non-competitive system, the higher functionality due to interdependency is not a result of coupling processes but is a result of the agents individually solving problems on a common domain and the diversity of their activity

interacting to provide higher system-wide functionality. This is similar to the last point, interdependency in selective systems is determined by the system dynamics, where interdependency in a non-selective system is expressed statically in the relationship between diversity and system functionality.

In a realistic system, both selective and non-competitive processes will coexist. The sources of diversity will be both from random and selective processes. The role of diversity will include both sources for selection and contributions to higher system functionality. What remains to be understood is the processes by which global system functionality encourages non-competitive traits on the individual level. As was observed in the simulations of group selection, the natural selection of cooperative traits disappeared when a single group became the largest unit. Evolution of cooperative traits requires closure to occur at a level greater than the individual groups.

Implications to closure

Revisiting the above arguments with respect to the need for closure in these systems leads to several conclusions. One is the role of emergence illustrated by both the simulations of group selection and of non-competitive problem solving. Emergence of properties that normally would be assumed in other models enables a closed system to exhibit sufficient functionality to be illustrative of realistic systems. The lesson is that if global restraints or assumptions are required to complete the model, then possibly some aspect of the model, e.g., spatial extent in the two example simulations above, is missing which would enable an emergent or natural closure of the system.

The second lesson is the possible importance of non-competitive mechanisms in higher system functionality as an alternative to selective mechanisms. Because the non-competitive mechanisms result in a direct coupling global system performance to the agents, it introduces a complex feedback mechanism, comparable to the interdependency of processes observed in the group selection. It is speculated that these non-competitive mechanisms are the missing pieces in the understanding of evolutionary systems and could be natural mechanisms for closure which have been missing.

References

- Johnson, N. L. (1998). *Collective Problem Solving: Functionality Beyond the Individual*, <http://ishi.lanl.gov/Documents1.html>.
- 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. <ftp://www.pscs.umich.edu/pub/papers/jb-bs-98.eps.gz>
- Prigogine, I. (1998). *The End of Certainty: Time, Chaos, and the New Laws of Nature*. N.Y.: The Free Press.
- Ridley, M. (1996). *Evolution*. (2nd edition). Cambridge, MA: Blackwell Science.
- Rothschild, M. (1990). *Bionomics: Economy as Ecosystem*. New York: Henry Holt.

The non-competitive simulations give a possible example of a trait that might have evolved due to a global selection, rather than from an individual selection. To the author's knowledge, most physical systems use a probabilistic sampling in either the governing equations or the characterization of their dynamics, whether the system is continuous or discrete. This is in contrast to the selection of a maximum signal in the non-competitive simulations. Even deterministic systems, such as planetary motion, are being treated as probabilistic systems to account for their chaotic nature (Prigogine, 1998). A molecular gas is an example of a discrete, self-organizing system, in which viscosity is an example of an emergent property. Turbulence in a macroscopic fluid is an example of a continuum system (here the emergent properties are the global structures and stochastics of the turbulence). In both of these systems, the ensemble average of many realizations is over a broad distribution, resulting from a probabilistic sampling of the possible space. Exceptions to the use of probabilistic sampling in physical systems are human creations, such as a transistor. To give an idea of how different physical problems might be with maximum sampling, a direct analogy to the non-selective simulations is the flow of water through a "maze" of pipes. If the flow system could select a maximum state instead of a probabilistic sample, water would only flow down the largest pipe and not down other pipes. Biological systems with sequential processes are in direct contrast to these physical systems in which the selection of a maximum state characterizes the selection process. The premiere examples are the working of the neurons in the brain or the selection of competing pheromone trails by an ant.

We speculate that the ability to select a maximum state in collective biological systems has evolved as a necessary capability to enable the functioning of a collective self-organizing system, which in turn has the desirable feature of being able to solve problems of greater complexity. Because the expression of the higher functionality is global, the adaption of this trait must be driven by global selection.

Acknowledgments Special appreciation of insightful conversations with Mark Bedau and Shareen Joshi. This work is supported by the Department of Energy under contract W-7405-ENG-36.