Applications of Self-Organization to Command, Control, and Coordination:

A Position Paper

Technical Report UT-CS-04-534

Bruce J. MacLennan[*]

Department of Computer Science
University of Tennessee, Knoxville
`www.cs.utk.edu/~mclennan`

November 17, 2004

Abstract

This position paper advocates the application of self-organization to large systems comprising robotic agents, informational (i.e. software) agents, and human beings. This would permit the construction of systems that are robust, flexible, and adaptive in a wide variety of hostile or dangerous environments, including those relevant to military operations, homeland security, and public safety.

---

1 Introduction

In this position paper we address two complementary problems. The first is how to use principles of self-organization to coordinate large numbers of independent agents (robots, autonomous vehicles, troops) to accomplish their missions in a way that is robust, flexible, and adaptive. The second problem is how to disrupt the command, control, and coordination structures of such self-organized systems when they are deployed by an enemy. To this end we advocate the investigation of general principles of self-organized command, control, and coordination that can be applied to a wide variety of militarily relevant systems of autonomous agents.

2 Rationale

2.1 Self-Organization in Natural Systems

2.1.1 Introduction

As is well known, many animal species exhibit robust, flexible, and adaptive command, control, and communication structures of operational relevance, which are implemented by distributed self-organization, rather than by centralized, hierarchical organization. It may be worthwhile to mention a few examples.

2.1.2 Army Ants

Army ant raids have been characterized as "the largest organized operations carried out regularly by any animal except humans" (Solé & Goodwin, 166). Although they have no fixed nest, colonies often comprise one million ants and may be up to 20 million in size. Several hundred thousand ants participate in the raids, which proceed from a temporary "bivouac" and may extend 350 meters. During the 15 days such a bivouac is occupied, a raid is launched in a direction that rotates 123° each day (thus avoiding over-exploitation of a region). Each such raid may cover 1000 m^2 of territory. After the vicinity of the bivouac has been thoroughly explored and exploited, the colony migrates for 20 days before establishing a new bivouac.

2.1.3 Ant Colony Self-Organization

Harvester ant colonies may have eight million members; yet, without any centralized control or information storage, such colonies are able to solve flexibly a number of important operational problems. The effectiveness of this collective decision-making has been compared to that of an individual vertebrate animal (Camazine & al., 245). Colonies continually allocate resources, such as adjusting the number of ants involved in foraging according to such factors as colony size, quantity of stored food, availability of food in the area, and the presence of competing colonies. They also solve optimization problems, such as finding the shortest paths to food sources, which they prioritize based on the quality of the food sources and on their distance and ease of access from the nest. These solutions are adaptive, in that if a path is blocked or a food source becomes unavailable, the paths will be reprioritized or new paths will be constructed.

Although a colony exhibits a 15-year life cycle, the workers live only a year, which shows that the colony's "memory" is not localized in individual workers. Rather it is stored in the physical structure of the nest and also distributed among the members of a continuously changing population of workers. Self-organization also occurs above the colony level; for example all the colonies of a particular species in a large geographic area (such as black ants in Southern England) coordinate the day of their "nuptial flight," by which they mate and establish new colonies.

2.1.4 Termite Colony Construction

Macrotermes termites construct large, elaborately and functionally structured mounds (Camazine & al., ch. 18), which may house a million individuals. The mounds may be 30m in diameter and 6 to 7m high (about 600 times the length of a worker); if termites were as large as people, these mounds would be a mile high and five miles in diameter. Since ventilation ports would permit the entry of predators, the mounds have radial ventilation fins, each containing a number of ducts to facilitate the diffusion of oxygen into the ducts and carbon dioxide out of them. These ducts also cause the air to circulate and regulate the temperature in the nest. The brood chambers and fungus gardens have a laminar structure supported by regular arrays of pillars. At least three pheromones are used by the termites, each of which may serve multiple functions. For example, the trail pheromone attracts workers to construction sites, encourages them to pick up soil granules (thus simultaneously clearing the trail and providing construction material), and regulates the size of the galleries to accommodate the level of traffic on the trail.

2.1.5 Wasp Nest Building

Wasps also construct nests with elaborate internal structures serving many functions. Simulations have shown that similar structures, as well as many other, potentially useful structures, may be assembled by autonomous agents obeying simple pattern-matching rules (Bonabeau et al., 1999). In some cases the result is completely deterministic, that is, it depends only on the rules (identical in all the wasps), not on the details of individual wasp behavior. In other cases, the structures may differ from run to run, although the qualitative structure (defined by macroscopic *order parameters*) is invariable. Some principles have been discovered by which multistage construction processes could be coordinated without the necessity of a project coordinator or an explicit behavioral program.

2.1.6 Flocks, Herds, and Schools

Groups of animals (flocks, herds, schools) can move like a single organism, executing quick maneuvers for more effective predation, to avoid predation, or for other purposes, such as hydrodynamic efficiency (Camazine & al., ch. 11). For example, a predator risks injury in attacking such a group, and a group can avoid a predator by evasive maneuvers, such as zigzagging or dividing into smaller groups. Conversely, some predators (tuna) coordinate their movement to surround prey, while others (killer whales) cooperate with each other, surrounding the prey (dolphin) and taking turns attacking. Also, in these highly coordinated maneuvers individuals rarely collide. In some cases very large schools of fish (millions of herring, hundreds of millions of cod) coordinate their behavior without leaders or global plans.

2.1.7 Slime Mold Self-Organization

The slime mold *Dictyostelium discoideum* has been a popular subject for investigations of self-organization and collective behavior (Kessin 2001). Although it might not be useful as a direct model for military applications, it does combine several principles of self-organization in a very sophisticated way. Although it is a single-celled organism (an amoeba), large numbers of individual amoebas respond to certain threats (e.g. depletion of food) in a highly coordinated fashion. The individuals assemble into disjoint “camps” of approximately the same size (50,000 to 100,000), and each group then reorganizes to behave like a multicellular organism, at some times like a migrating slug, at others like a plant for dispersing spores. In the process we find large numbers of essentially identical cells self-organizing into functionally distinct groups constituting the parts of the slug. This self-organization is coordinated by a sophisticated distributed signaling system and by highly coordinated cell movement (both before and after assembly of the multicellular slug). The slug has no brain or nervous system, yet it is able to discriminate minute temperature and light gradients and migrate in a more favorable direction. If a slug is severed or partially eaten (by a nematode), then the cells in the remainder reassign themselves to various functions so that the remainder becomes a functionally complete slug.

2.1.8 Morphogenesis

Initially identical cells in an embryo differentiate and self-organize into a variety of tissues and organ structures with a complex microstructure and macrostructure. Many tissues can repair themselves throughout the organism’s life, and in some species complex structures (e.g. limbs) can be regenerated. Organisms provide many examples of the use of self-organizational principles to maintain operational effectiveness in spite of damage, disease, and other impediments.

2.2 Self-Organization for Artificial Systems

2.2.1 Potential Domains of Application

The principles of self-organization, once well understood, have many applications in warfare, security, and safety; indeed, they could revolutionize operations in these areas (Adams 2000). Before discussing some of the potential benefits of self-organization, it will be worthwhile to mention some of the systems to which they may be applied. We refer to these human-designed systems as *artificial systems* in order to distinguish them from the natural systems we have been discussing.

2.2.1.1 Robots and Autonomous Vehicles

The most obvious application of the principles of self-organization is to various kinds of robots and autonomous vehicles deployed on the land or sea, or in air, space, or undersea. We have in mind both contemporary robotic systems and future microrobotic agents and nanobots (including genetically-engineered microorganisms). In the latter cases especially we can anticipate that the agents will be inexpensive enough to be deployable in large enough numbers to make self-organization a critical enabling technology. As will be discussed in more detail later, such artificial agents have many of the same operational needs as the natural agents mentioned above.

2.2.1.2 Informational Agents

Many of these same principles of self-organization can be applied to non-physical (i.e. software) *informational agents* operating on the internet or in similar information environments. Under some conditions *memes* (units of meaning) can operate as informational agents (analogous to the analysis of "selfish genes" as intentional agents).

2.2.1.3 Command, Control, and Coordination of Human Agents

Self-organized systems in nature demonstrate how relatively unintelligent agents (such as ants) can cooperate and achieve higher levels of collective intelligence (comparable to that of a vertebrate animal). This, of course, is part of what makes self-organization attractive for coordinating microrobotic agents and small autonomous vehicles. However, such principles also operate in the interactions of intelligent agents, such as human beings and other complex animals, and so the principles of self-organization may be applied to command, control, coordination, and communication operations for troops and other human teams. For example, individual warfighters or small units may carry compact devices facilitating valuable functions such as robust communication, distributed information storage, task assignment, distributed intelligence gathering, and coordinated maneuvering or other action. Self-organization may be an advance on current concepts of distributed operations (Schmidle & Hoffman 2004).

In the following paragraphs we will highlight some specific applications of self-organization to artificial systems comprising robotic, informational, and/or human agents.

2.2.2 Allocation of Resources

As in ant colonies and many other natural systems, self-organized resource allocation is effective, efficient, flexible, and adaptive. An especially important resource allocation problem is the detailing of agents (human or otherwise) to particular tasks, in a manner that responds agilely to changing conditions and availability of assets. Self-organization accomplishes this function without the necessity (and vulnerability) of centralized information gathering and decision-making. Related operations include the stockpiling, dispersal, and relocation of materiel and analogous redistribution of information.

2.2.3 Exploration vs. Exploitation

Related to resource allocation is the relative distribution of effort to *exploration* vs. *exploitation*. In many situations it is difficult to say when enough information has been gathered in order to take appropriate action; we run the risk of either waiting until we're sure, when it may be too late, or of acting inappropriately on the basis of insufficient information. Many self-organized systems avoid this dilemma by gradually shifting effort from relatively unbiased exploration toward increasingly directed exploration and exploitation. These approaches may achieve near-optimal relative allocation of resources to exploration and exploitation (e.g., Holland's analysis of the k-armed bandit problem). Such systems are also adaptive in that they automatically move back toward less directed exploration as it becomes apparent that initial information was misleading or as conditions change.

2.2.4 Communication

Many natural self-organized systems function effectively without long-distance directed or broadcast means of communication. Instead, strictly local interactions serve to coordinate global behavior. Furthermore, global information is estimated statistically through local interactions, rather than being gathered, forwarded to centralized repositories, and redistributed throughout the system. Distributed communication by means of local interactions is more reliable because it is less sensitive to localized disruptions (or embedded hostile agents) and less susceptible to interception. Also, since it typically operates over relatively short distances, it is less subject to jamming, and since the communication is among pairs or small groups of agents, peer-to-peer rather than global encryption systems can be used, thus improving security.

2.2.5 Distributed Synchronization

Many operations must be coordinated and synchronized for maximum effectiveness, but centralized synchronization runs the risks of interception, counterfeiting, or jamming; like any centralized mechanism, it is an Achilles heel. An alternative, distributed approach is found in nature, for example in the synchronized nuptial flights of black ants and in the synchronized flashing of some species of fireflies. Similar techniques can be employed for self-organized synchronization of artificial agents, with corresponding improvements in reliability and security.

2.2.6 Information Storage

Self-organized systems in nature typically make use of distributed information storage, but this is not simply a matter of having explicit information replicated at multiple locations. Rather, we find global information to be *implicit* in the limited memories of the agents and in the artifacts constructed by them (e.g. pheromone trails, nest structure). Aspects of this information are reconstructed locally when needed to guide the behavior of individual agents. As a consequence the information base is robust; eliminating one or a few centralized repositories cannot destroy it, and the quality of the information retrieved degrades gracefully with increasing damage to the total system. Distributed, implicit storage also decreases vulnerability to corruption of the information base or to intentional incorporation of misinformation by hostile agents.

2.2.7 Construction

Just as social insects self-organize to build complex nests, robots can self-organize to construct complex structures (e.g., roads, barriers) or to relocate objects or materials (e.g., mine clearing, hazardous material removal). Similarly, self-organization can guide robotic or human agents to assemble large spatial structures (battle lines, supply routes, etc.), and informational agents can construct emergent information structures. Self-organized approaches to construction are robust, flexible, and adaptive; they succeed in the absence of centralized control, detailed blueprints, or explicit procedures, all potential vulnerabilities.

2.2.8 Summary

These are just a few examples of the application of self-organization to command, control, and coordination. Consideration of the natural models will suggest many other applications of operational relevance.

3 Approach

If our goal is to identify broad principles of self-organization that can be put to use in a wide variety of artificial systems comprising robotic, informational, and human agents, then we should begin with a two-pronged approach comprising a literature review and simulation studies.

3.1 Literature Review

In recent decades there has been a great deal of research in self-organized natural systems, some of which has been gathered into valuable surveys, such as Camzine et al. (2001). Much of this research has been conducted by biologists and other scientists more interested in understanding natural self-organized systems than in the application of self-organization to practical problems of command, control, and coordination in warfare, homeland security, and public safety operations. Therefore, the first step is to mine this literature for general principles that can be applied effectively to these problems. To illustrate the sorts of principles that we seek, we will mention a few that have been identified already.

3.1.1 Examples of Principles

3.1.1.1 Positive and Negative Feedback

Positive and negative feedback are familiar ideas from engineering, but they have a special role in self-organization. For example, positive feedback biases exploration into directions that have already proved valuable, so that the system can begin exploiting information before it has finished gathering it, and negative feedback limits commitment to particular options, so that the system remain flexible and adaptive. Their interaction can achieve a near-optimal, adaptive balance between exploration and exploitation. Also, positive and negative feedback processes often work in consort to create determinate structures extended in space and / or time; positive feedback creates and increases the structures, while negative feedback limits and differentiates them. The combination of positive and negative feedback appears in many forms in self-organized systems: growth / limit, amplification / stabilization, cooperation / competition, etc.

3.1.1.2 Diffusion

Many self-organized systems in nature make productive use of one of the simplest of physical processes: *diffusion*. For purposes of self-organization, diffusion may be implemented by the physical diffusion of chemical substances or by the random locomotion of agents (robots, nanobots, microorganisms, informational agents); troops may even be guided to engage in non-directed interactions that amount to diffusion. Diffusion may be used for many purposes in self-organization. First, it may be used as a robust means of broadcasting information. Second, it may be used as a mechanism for unbiased parallel search, since each individual diffusing agent is exploring the space in which it is wandering; other self-organizational mechanisms may exploit this distributed, implicit information. Diffusion also makes agents available for tasks determined by local conditions without centralized task assignment. Finally, differential diffusion rates and decay rates can be used (via positive and negative feedback) to create spatial structures and to control the balance between exploration and exploitation in optimization processes. In summary, diffusion is one mechanism by which we can make productive use of very large numbers of behaviorally simple agents.

3.1.1.3 Noise, Randomness, and Imperfection

When large numbers of imperfect agents operate in a complex, noisy, unpredictable, real-world environment, and interact with other imperfect, and perhaps hostile agents, we must expect a great deal of noise, randomness, and imperfection in all behavioral processes. On the one hand, self-organized systems, with their built-in redundancy and distributed organization, are well suited to deal with these problems. On the other, they often make productive use of randomness, noise and imperfection as part of their operation. For example, random interactions can be used for statistical estimation of relevant global variables, and imperfect path following leads to exploration, flexibility, and adaptability in foraging. As in simulated annealing and related optimization algorithms, a certain number of locally “wrong” decisions can help a system escape from a suboptimal state and approach a more nearly optimal solution to a problem. Also, we have already mentioned diffusion, which depends on random walks. Therefore, the *productive use of randomness, noise, uncertainty, and imperfection* is an important principle of self-organization.

3.1.1.4 Amplification of Random Fluctuations

Many self-organizational processes begin with the *amplification* (through positive feedback) *of initial random fluctuations*. This breaks the symmetry of the initial state, but often in unpredictable but operationally equivalent ways. That is, the job gets done, but hostile forces will have difficulty predicting precisely how it gets done. For example, agents may self-organize into groups of a predictable size and average distance between groups, but the exact location of the groups will be unpredictable.

3.1.1.5 Stigmergy

Stigmergy is an important principle of self-organization, seen for example in wasp nest building and spider web construction. It refers to a way of coordinating a collective construction process so that the project itself contains the information necessary to guide the actions of the workers. Therefore there is no need for an external blueprint or project leader (both potential vulnerabilities), nor do the workers need an internal explicit program for the entire project (leading to inflexible group behavior). In effect, the evolving project serves as a medium of information storage, communication, and coordination. As a consequence, in many cases agents can adapt automatically to disruptions in the construction process, such as the destruction of some or all of the structure; they respond appropriately to what remains. Each agent encountering the project knows from the project itself what needs doing. Stigmergy permits the use of simpler agents and decreases direct communication between agents. It also permits greater flexibility in assignment of agents to tasks, since the task-relevant information is stored in the project rather than in the agents. There are two different kinds of stigmergy (Bonabeau & al., 205–8). In *continuous* or *quantitative* stigmergy, quantitatively different stimuli trigger quantitatively different behaviors (often subject to positive and negative feedback). In *discrete* or *qualitative* stigmergy, stimuli are classified into distinct categories, which trigger distinct behaviors. Biologists have discovered a number of supporting principles that improve the effectiveness of stigmergy in complex construction processes; we could apply these to self-organized command, control, and coordination of operations.

3.1.1.6 Simple Local Microdecisions

One of the central principles of self-organization is that *complex, adaptive macrobehavior emerges from simple, local microdecisions.* This principle is what allows simple agents (or complex agents obeying simple rules) to collectively solve difficult problems in real time. Indeed, having agents following complex behavioral strategies may be no advantage, since simple strategies are often more robust and flexible. Agents' microdecisions often are based on local patterns in signals (e.g. local concentration gradients or concentration variations over certain periods of time) and between signals of different types. Agents are also aware of neighboring agents and their behavior, and can communicate with them via simple signals. As already mentioned, such local, short-range communications are resistant to interception and jamming.

3.1.1.7 Multiple Interactions

Most self-organized systems depend on *multiple interactions* among the agents, which allow information to propagate among the agents and the emergence of long-distance coordination. To facilitate these interactions agents may move around, randomly as well as in accord with their microdecisions.

3.1.1.8 Circular Causality

A well-known principle of self-organization is *circular causality* (also known as *the macro/micro feedback loop* or *collectivism*). That is, global order emerges from the interaction of the agents, which in turn respond to the global order. Although we can see how it operates in many natural systems, using it in artificial systems to achieve specific ends is more difficult, precisely because we are dealing with complex systems, and the global behavior is emergent, and therefore difficult to predict. Thus, one goal of our research is a better understanding of this principle.

3.1.1.9 Excitable Media

In an *excitable medium*, local relaying of information can be used to construct large static or moving patterns or structures, which are qualitatively predictable (e.g., in size), so they can serve some purpose, but are unpredictable in detail (e.g., exact location), so they are more resistant to preemptive hostile action. Examples of excitable media include cardiac tissue, neural cortex, aggregating groups of slime molds, the Belousov-Zhabotinski reaction, and the spread of certain infectious agents. We should investigate how to apply this principle so that agents can organize themselves in space and time.

3.1.1.10 Flexible, Adaptive Stationary States

Many of the structures created by self-organizing systems, such as ant paths, are *stationary states* of a dynamical system. That is, they are stable only insofar as the conditions under which they were created are stable. When the conditions change, the equilibrium automatically adjusts itself to a new stationary state, so the system is adaptive. Similarly, if the structure is damaged in some way, the system will re-establish the stationary state, thus implementing self-repair.

3.1.1.11 Nonconvergence, Diversity, and Suboptimal Solutions

In traditional optimization, we typically want our algorithms to converge to a unique solution. Natural self-organized systems, in contrast, frequently settle on several good, but non-optimal solutions. *Non-convergence*, an apparent limitation, is often advantageous, however, since having a diversity of good solutions allows the system to be more robust and adaptive.

3.1.1.12 Entrainment and Distributed Synchronization

Various mechanisms are used by natural systems to achieve self-organized, distributed synchronization. One of these is the *entrainment* of coupled oscillators, but a complete analysis of large coupled systems is still lacking (Winfree 2001).

3.2 Simulations

Any study of self-organized systems, especially self-organized complex adaptive systems, must depend on computerized simulations, for our analytical tools in this area are very limited. Indeed, the wide availability of inexpensive, powerful computers has enabled much of the recent progress in complex systems and self-organization. Therefore, after the identification of candidate principles of self-organization from the literature, the next step should be to use simulation tools to investigate the application of these principles to various problems of operational significance. The first simulations should be comparatively simple and abstract in order to explore the principles in their simplest context and to demonstrate their operation in the most general terms (i.e. so they are not restricted to particular operational contexts). Later simulations can be more detailed and specific, exploring self-organization in simulated systems of operational relevance. It is impossible to say at this time which systems should be simulated, but they should include artificial systems comprising all the kinds of agents (robotic, informational, human), that is, they should be operationally more realistic simulations oriented to problems relevant to military operations, homeland security, and public safety.

4 Potential Value

4.1 Robust, Adaptive Command, Control, and Coordination of Friendly Systems

The primary value of the research we advocate is to identify principles of self-organization that can be applied to the robust and adaptive implementation of command, control, and coordination for large artificial systems of agents (robotic, informational, or human). Such systems will be increasingly important for homeland security, military operations, and public safety applications.

4.2 Disruption of Hostile Self-Organized Systems

While the principles of self-organization offer may benefits to friendly forces, we must assume that they will also be applied by our opponents (as they already are by terrorist organizations, guerillas, etc.). Therefore, complementary to valuable principles for constructing self-organized friendly systems, we should investigate means for interfering with the self-organization of hostile systems. Here too we have much to learn from nature, for many species are in close competition, and we can apply our knowledge to hostile agents, both human and nonhuman (e.g., pathogens, robots). That is, knowledge of the principles of self-organization is critical to discovering and defeating self-organized systems, as well as to constructing and deploying them.

5 Summary

In summary, scientific research has shown the pervasiveness of self-organization in the natural world, from nonliving systems, through microorganisms, to species of all degrees of complexity, including human beings. This research has demonstrated how comparatively simple interactions, often among organisms with limited cognitive capacities, can solve complex command, control, and coordination problems in order to promote their survival and to accomplish their ends. The behaviors of these species is more robust, flexible, and adaptive than they would be if they were not based on self-organization. With this increased knowledge of natural self-organization, has come improved understanding of various general principles that can be applied to artificial systems to achieve the same benefits. In past research, a variety of simulation studies have shown that these principles of self-organization can be applied in artificial systems, which may be quite different from the natural systems in which the principles were originally observed. Self-organization is especially attractive as an approach to the robust, flexible, and adaptive implementation of command, control, and communication systems of enormous potential value to military operations, homeland security, and public safety.

6 Bibliography

Adams, T. K. (2000). The Real Military Revolution. *Parameters*, Autumn 2000: 54–65.

Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems.* Oxford.

Camazine, S., Deneubourg, J.-L., Franks, N. R., Sneyd, J., Theraulaz, G., & Bonabeau, E. (2001). *Self-Organization in Biological Systems.* Princeton.

Gordon, D. (1999). *Ants at Work.* Free Press.

Ilachinski, A. (1996). *Land Warfare and Complexity, Part I: Mathematical Background and Technical Sourcebook (U).* Center for Naval Analyses.

Kessin, R. H. (2001). *Dictyostelium: Evolution, Cell Biology, and the Development of Multicellularity.* Cambridge.

Schmidle, R. E., & Hoffman, F. G. (2004). Commanding the Contested Zones. *U.S. Naval Institute Proceedings*, September 2004.

Solé, R., & Goodwin, B. (2000). *Signs of Life: How Complexity Pervades Biology.* Basic Books.

Winfree, A. T. (2001). *The Geometry of Biological Time*, 2nd ed. Springer.