Presented under the *Modeling and Simulation* and *Network-Centric Application* Topics
for the 8th International Command and Control Research and Technology Symposium,
June 17 – 19, 2003, National Defense University Washington, DC

Paper Title: ***Agent-Based Modeling for Testing and Designing Novel Decentralized Command and Control System Paradigms***

Co-Author's Name: Dr. Eric Bonabeau, Ph.D., Icosystem Corporation
Title: Chief Scientist, Icosystem Corporation
Address: 10 Fawcett Street, Cambridge, MA 02138; **Phone:** 617-520-1020
Fax: 617-492-1505; **Email:** eric@icosystem.com

Co-Author's Name: LTC Carl W. Hunt, Ph.D., United States Army
Current Assignment: Student, US National War College, Ft. McNair, Washington, DC
Address: 8652 Morning Star Ct, Springfield, VA 22153; **Phone:** 703-451-0604
Email: huntc@ndu.edu / carl.hunt@us.army.mil

Co-Author's Name: Dr. Paolo Gaudiano, Ph.D., Icosystem Corporation
Title: Senior Scientist, Icosystem Corporation
Address: 10 Fawcett Street, Cambridge, MA 02138; **Phone:** 617-520-1020
Fax: 617-492-1505; **Email:** paolo@icosystem.com

Agent-Based Modeling for Testing and Designing Novel Decentralized Command and Control System Paradigms

Eric Bonabeau
Icosystem Corporation, 10 Fawcett Street, Cambridge, MA 02138; Email: eric@icosystem.com

LTC Carl W. Hunt
United States Army, US National War College, Ft. McNair, Washington, DC; Email: huntc@ndu.edu / carl.hunt@us.army.mil

Paolo Gaudiano
Icosystem Corporation, 10 Fawcett Street, Cambridge, MA 02138; Email: paolo@icosystem.com

Abstract

Agent-based modeling (ABM) is a recent simulation modeling technique that consists of modeling a system from the bottom up, capturing the interactions taking place between the system's constituent units. Such a bottom up approach enables users to describe and predict emergent phenomena. These include aggregate, system-level behaviors that can be counter-intuitive. Because decentralized command and control (DC2) paradigms can sometimes lead to counter-intuitive phenomena, ABM is the tool of choice to test DC2 and can provide significant insight into the design of DC2 approaches.

1. Decentralized Command and Control: Benefits and Challenges

Testing new decentralized command and control (C2) paradigms can be challenging as decentralization requires a fresh mindset. New C2 methods can lead to unexpected, unanticipated system-wide behavior – also called emergent phenomena. While decentralization and distributed organizations offer clear benefits over traditional C2 approaches (robustness, flexibility, fluidity, responsiveness, self-organization), they also present clear challenges when it comes to designing the behavioral rules that, for example, individual soldiers must follow on the battlefield.

For example, how does a commander test the rules of engagement to make sure that the system (ranging from platoon to brigade, or from swarms of unmanned aerial vehicles (UAV) to fleets of warships) can achieve its mission, meet specific requirements and not break down under rare but not impossible pathological conditions? Only when reliable testing tools become available will decentralized C2 paradigms be adopted. Such tools are necessary to reach the required level of confidence in new approaches and shift from a centralized C2 mindset to a decentralized one.

An important solution to this dilemma has surfaced in the last few years: agent-based modeling (ABM) (Bonabeau, 2000, 2002a, 2002b; Casti, 1997; Epstein & Axtell, 1996; Hunt, 2001). In ABM, systems are modeled as collections of autonomous decision-making entities, called agents. Each agent individually assesses its situation and makes decisions based upon a set of rules. Agents may execute various behaviors appropriate for the system they represent. These systems, for example, may include searching, attacking, or performing battle damage assessment for an agent-based model of sensors such as UAVs. Repetitive, competitive and cooperative interactions between agents are a feature of agent-based modeling, where such modeling relies on the power of computers to explore dynamics out of the reach of pure mathematical methods.

At the simplest level, an agent-based model consists of a system of agents and the relationships between them. Even a simple agent-based model can exhibit complex behavior patterns and provide valuable insights about the dynamics of the real-world system that it emulates. In addition, agents may be capable of evolving, thereby allowing unanticipated behaviors to emerge. Sophisticated ABM sometimes incorporates neural networks and genetic algorithms to allow realistic learning and adaptation.

In recent years, ABM has been successfully applied to a range of problems in the commercial and civilian world (supply chain modeling, crowd modeling for evacuation, flow management in public spaces, theme parks and supermarkets, traffic control, organizational modeling, risk management, and market modeling) as well as in the military and law enforcement world (battlefield simulation, UAV control, military communications network security, and drug trafficking, for example). In the rest of this paper we will present the fundamental principles of ABM, present some applications of ABM relevant to command and control, and introduce ways of designing new DC2 paradigms using ABM as a testing platform.

2. Emergent Phenomena and Agent-based Modeling

Emergent phenomena result from the interactions of individual entities. By definition, they cannot be reduced to the system's parts: the whole is more than the sum of its parts because of interactions among the parts. An emergent phenomenon can have properties that are decoupled from the properties of the part. For example, a traffic jam, which results from the behavior of, and interactions between individual vehicle drivers, may be moving in the direction opposite to that of the cars that cause it. This characteristic of emergent phenomena makes them difficult to understand and predict: emergent phenomena can be counterintuitive.

We will review a variety of examples of counterintuitive emergent phenomena in the following sections. Agent-based simulation is by its very nature the canonical approach to modeling emergent phenomena: in an agent-based simulation, one models and simulates the behavior of the system's constituent units (the agents) and their interactions, capturing emergence when the simulation is run. Perhaps the simplest illustration of emergence is Boids, Craig Reynolds' virtual creatures (Reynolds, 1987). Boids move in flocks (think of a flock of boids as a simple model of a flock of birds or a school of fish), they obey four simple rules, which characterize their behavior entirely:

- Separation: steer to avoid crowding local flockmates
- Alignment: steer towards the average heading of local flockmates
- Cohesion: steer to move toward the average position of local flock-mates
- Avoidance: steer to move away from oncoming obstacles

When simulated in a computer, these four simple rules lead to an appearance of global coordinated behavior strongly reminiscent of real flocks or schools. For example, boids collectively avoid obstacles, a property that is not explicitly programmed into their behavior. In other words, collective obstacle avoidance is an emergent property.

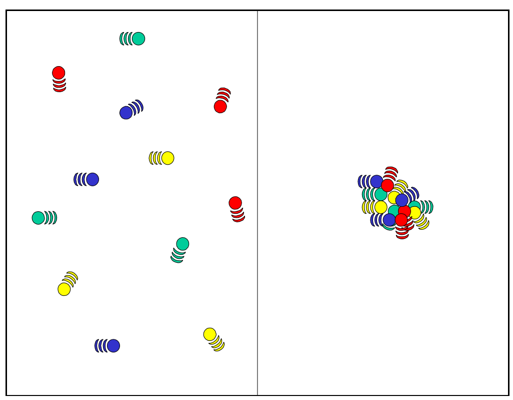

Figure 1

The Aggressor-Protector Game. Left: when everyone moves to keep A in between themselves and B, the group seems to be moving chaotically, covering the whole space. Right: when the rule is modified slightly, the group collapses onto itself, forming one compact cluster.

Figure 1 above depicts another example of an emergent phenomenon involving people (Bonabeau, 2002a, b). It starts with a game that is easy to play with a group of 10 to 40 people. Ask an audience to each randomly select 2 individuals – person A and person B. Now ask them to move so that they always keep A in between themselves and B, so A is their Protector from B. Everyone in the room will mill about in a seemingly random fashion and will soon begin to ask why they are doing this. At this point, tell them that they are now the Protector and ask them to move so that they keep themselves in between A and B. The results are striking: almost instantaneously the whole room will implode on itself with everyone clustering together in a tight knot. This example shows how simple individual rules can lead to coherent group behavior and how small changes in those rules can have a dramatic impact on the group behavior. This also shows how intuition can be a very poor guide to outcomes beyond a very limited level of complexity. The group's collective behavior is an emergent phenomenon.

By using a simple agent-based simulation (available at: www.icosystem.net/game.html) in which each person is modeled as an autonomous agent following the rules, one can actually predict the emerging collective behavior. Although this is a simple example, where individual behavior does not change over time, ABM enables one to deal with more complex individual behavior, including learning and adaptation.

One may also want to use ABM when there is the potential for emergent phenomena, i.e., when:

- Individual behavior is nonlinear, can be characterized by thresholds, if-then rules or nonlinear coupling. Describing discontinuity in individual behavior is difficult with differential equations.

- Individual behavior exhibits memory, path-dependence and hysteresis, non-markovian behavior, or temporal correlations, including learning and adaptation.

- Agent interactions are heterogeneous and can generate network effects. Aggregate flow equations usually assume global, homogeneous mixing but the topology of the interaction network can lead to significant deviations from predicted aggregate behavior.

- Averages won't do the job. Aggregate differential equations tend to smooth out fluctuations; agent-based modeling typically does not. This is important because under certain conditions, fluctuations can be amplified – the system is linearly stable but can be nonlinearly unstable to even small perturbations

Most often, ABM is also the most natural for describing and simulating a system comprised of "behavioral" entities. Whether one is attempting to describe a stampede, a traffic jam, the stock market, voters or how an organization works, agent-based modeling makes the model look closer to reality. For example, it is more natural to describe how shoppers move in a supermarket than to come up with the equations that govern the dynamics of the density of shoppers. Because the density equations result from the behavior of shoppers, the agent-based modeling approach will also enable the user to study aggregate properties.

Agent-based modeling also makes it possible to realize the full potential of the data a company may have about its customers: panel data and customer surveys provide information about what actual people actually do. Knowing the actual shopping basket of a customer makes it possible to create a virtual agent with that shopping basket rather than a density of people with a synthetic shopping basket computed from averaging over shopping data.

One may want to use ABM when describing the system from the perspective of its constituent units' activities as more natural; i.e., when:

- The behavior of individuals cannot be clearly defined through aggregate transition rates.

- Individual behavior is complex. Everything can be done with equations, in principle, but the complexity of differential equations increases exponentially as the complexity of behavior increases. Describing complex individual behavior with equations becomes intractable.

- Activities are a more natural way of describing the system than processes.

- Validation and calibration of the model through expert judgment is crucial. ABM is often the most appropriate way of describing what is actually happening in the real world, and the experts can easily "connect" to the model and have a feeling of "ownership".

- Stochasticity applies to the agents' behavior. With ABM, sources of randomness may be applied to the right places as opposed to a noise term added more or less arbitrarily to an aggregate equation.

Another use of ABM that suggests ways to enhance the process of inference and discovery in intelligence and criminal investigation situations is found in the Agent Based Evidence Marshaling (ABEM) model (Hunt, 2001). The ABEM model demonstrates how rich and complex hypotheses can grow and be tested based on the simple passing of very basic information sets called tuples. A simplified rule set accommodates the movement of tuple-encoded intelligence and criminal investigative evidence to pass within a visual environment in ways to depict hypothesis formation and testing. Out of the interactions empowered in the ABEM environment, agent discovery and learning take place, prompting the human analyst to infer or discover new lines of inquiry that could be pursued. The ABEM model has recently been extended to other intelligence problems such as those encountered in the imagery collection and analysis world.

3. A New Mindset

ABM is a mindset more than it is a technology. The ABM mindset consists of describing a system from the perspective of its constituent units. The benefits of ABM over other modeling techniques can be presented in two statements: (1) ABM captures emergent phenomena; and (2), ABM provides a natural description of a system. Emergent phenomena result from the interactions of individual entities. By definition, they cannot be reduced to the system's parts: the whole is more than the sum of its parts because of interactions among the parts.

An emergent phenomenon can have properties that are decoupled from the properties of the part. For example, a traffic jam, which results from the behavior of, and interactions between individual vehicle drivers, may be moving in the direction opposite to that of the cars that cause it. This characteristic of emergent phenomena makes them difficult to understand and predict as emergent phenomena can be counterintuitive to an understanding of constituent components.

ABM is by its very nature the canonical approach to modeling emergent phenomena. In agent-based simulations, a commander may model and simulate the behavior of the system's constituent units (the agents) and their interactions, capturing emergence when the simulation is run. Being able to predict the collective behavior of the group that emerges from individual behavior is important because it enables commanders and their staffs to test a wide range of individual and collective behavioral rules and to observe if and how robustly the rules lead to the desired collective-level outcome.

Recent developments in ABM clearly indicate that ABM provides a natural and visual description of the system as opposed to, for example, density equations of motion. Because ABM replicates the way the real agents behave *in silico*, it is straightforward to integrate specialist knowledge and expertise into it. Domain experts can relate to the model because it replicates the way they see the world. Today's models may be designed so they are simple enough that informed non-experts can glean relevant insights. This desirable property makes ABM a compelling proposition to promote the adoption of novel DC2 paradigms.

4. Distributed Command and Control (DC2) Management of Flows and Traffic

Managing flows and traffic can be extremely important for the military. For example, designing safe public infrastructures or evacuating a large, densely populated area are two problems that seem to demand DC2 approaches. This is in large part because centralized control is often brittle or too single-focused and thus can often fail to respond to issues that arise in real time. Fortunately, robust distributed C2 approaches can be tested using ABM.

Scientists in academia and business have built flow models for years in attempts to understand, predict and possibly control all kinds of flows. Typical flow models, based on equations that describe how the density of entities (such as people, vehicles or other agents) vary with time are suitable when the behavior of the agents can be described in simple terms and when averages are appropriate; that is, when the behavior of individual agents is relatively smooth. But when it comes to describing how a person or even worse a group of people might respond to a fire alarm in a movie theatre, the behavior of each person is very far from being smooth and different people may have very different behaviors. In this case, ABM is the best approach to modeling panic behavior. Applications include emergency evacuation planning, flow optimization, and waiting time minimization.

For example, crowd stampedes induced by panic often lead to fatalities as people are crushed or trampled. Such phenomena may be triggered in life-threatening situations such as fires in crowded buildings or may arise from the rush for seats or sometimes seemingly without causes. Recent tragic examples include the panics in Harare, Zimbabwe, and at the Roskilde rock concert in Denmark.

Although engineers are finding ways to reduce the magnitude of such disasters, their frequency seems to be increasing, as growing population densities combined with easier transportation lead to greater mass events such as pop concerts, sporting events, and demonstrations. These kinds of circumstances are likely to continue to grow. Panicking people are obsessed by short-term personal interests uncontrolled by social and cultural constraints. This is possibly a result of reduced attention in situations of fear, thus it is possible that in such cases alternatives like side exits seem to be mostly ignored.

In addition, there is social contagion; that is, a transition from individual to mass psychology, in which individuals transfer control over their actions to others leading to conformity. Such irrational herding behavior often leads to bad overall results like dangerous overcrowding and slower escape, increasing the fatalities or, more generally, the damage. In agent terms, collective panic behavior is an emergent phenomenon that results from relatively complex individual-level behavior and interactions between individuals (hypnotic effect, mutual excitation of a primordial instinct, circular reactions, social facilitation).

ABM is ideally suited to provide valuable insights into the mechanisms of and preconditions for panic and jamming by in-coordination. Simulation results (Helbing et al. 2000; Still 1993, 2000) suggest practical ways of minimizing the harmful consequences of such events and the existence of an optimal escape strategy, corresponding to a suitable mixture of individualistic and collective behavior. For example, let's consider a fire escape situation in a confined space such as a movie theatre, a concert hall or the Superbowl stadium. For the purpose of this high-level discussion, assume that there is one exit available.

First, what would be the optimal strategy for getting out as fast as possible? Most people intuit that crowds act in the same way as grains of sand in an egg timer, where the fastest flow is down the central axis. Crowds are not affected by friction or gravity in the same way as grains of sand. At the edges of the queue a member of a crowd is aligned with the flow profile and has fewer interactions. Movement towards the open doorway is less impeded and consequently faster.

Another question, which has to do with the design of a facility, is: how can one increase the outflow of people? Narrowing down the problem a bit, one could ask: what is the effect of putting a column (a pillar) just before exit, slightly asymmetrically (for example, to the exit's left), about one meter away from the exit? Intuitively, one might think that the column will slow down the outflow of people. However, agent-based simulation, backed by real-world experiments, indicates that the column regulates the flow, leading to fewer injured people and a significant increase in the flow, especially if one assumes that injured people can't move and impede the flow (Helbing et al. 2000).

As Figure 2, below, demonstrates, the insertion of a column to separate escaping people tends to increase the number of people who can successfully extricate themselves from the dangers brought on by stampede. While it may not seem intuitive that this would be the case, this simple agent-based model and the accompanying table data show that more than 70 agents may escape uninjured in 45 seconds as opposed to nearly 45 who may escape from the same environment without the intervening column. ABMs such as those described in Helbing can assist in modeling and discovering non-intuitive findings.

Type of Simulation	Escaped until t=45s	Injured until t=45s
Without column, injured people don't move (Stampede / 200 people)	44	5
With column, injured people don't move (Column)	72	0

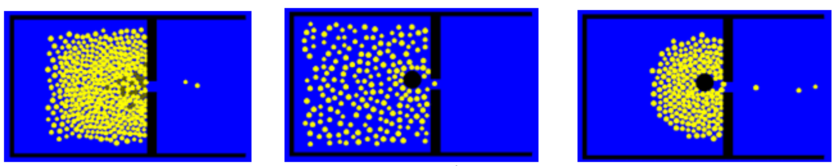

a b c

Figure 2
Fire escape agent-based simulation (live simulation available at www.helbing.org). People are represented by circles, green (darker) circles being injured people. (a) No column. (b) With column, after 10 seconds. (c) With column, after 20 seconds. After Helbing et al. (2000).

a. DC2 for the supply chain

DC2 strategies may be used to organize a supply chain in a robust manner so it can re-organize itself in response to glitches and perturbations. An interesting illustration of a DC2 strategy in this context comes from Southwest Airlines' cargo activities (Bonabeau & Meyer, 2001). In 1998, Southwest Airlines was having trouble with cargo bottlenecks that delayed the arrival of packages and increased operating costs. At the time, employees were trying to load freight onto the first plane going in the right direction—a seemingly reasonable strategy. But because of it, workers often had to offload cargo from a flight and store it—sometimes overnight—until the appropriate aircraft was available.

Researchers studying this problem then discovered something surprising: it is often better to leave cargo on a plane even though the vehicle might be headed in the wrong direction. In other words, a circuitous routing can, all things considered, actually be more efficient. The new approach has cut the cost of freight-handling labor by 20%. In addition, fewer planes are now flying full, which opens up significant opportunities for the company to generate new business. Thanks to the improvements, Southwest estimated an annual gain of more than $10 million. One of the interesting characteristics of the Southwest solution is its reliance on simple rules that ramp personnel can easily interpret and make sense of. Southwest Airlines wanted to keep cargo routing “low tech” to make it more adaptable: ramp personnel should not have to use a complex information system to make simple package transfer decisions.

Using an agent-based model, Southwest Airlines was able to test sets of simple rules and select the ones that maximized the efficiency of the operations and are robust to such perturbations as delays, bad weather or even fluctuations in demand. ABM was particularly useful here because the impact of implementing a particular set of simple local routing rules can sometimes be felt throughout the entire system and is often counter-intuitive and/or unpredictable. For example, the policy that was implemented prior to the ABM project often led to bottlenecks at hubs because the shortest path between a given airport and another airport often goes through a hub; accepting longer routes decreases congestion at the hubs.

Without an ABM to explore such emergent properties of the system, it may have been impossible for Southwest Airlines to discover what the problem really was and how to solve it. The resulting set of routing rules, in addition to significantly improving the cargo activities’ efficiency, also gave the whole system more flexibility to external perturbations (weather, delays, fluctuations in demand), and robustness (mechanical problems, flight cancellations) by empowering local personnel to make decentralized decisions. This example is obviously applicable to many logistics and supply chain situations in the military world: ABM can be used to test and design DC2 logistics management rules.

b. DC2 for the organization

Using ABM, it is possible to design an organization that performs effectively under a DC2 management strategy. A useful illustration of this is inspired by task allocation in a biological organization (Bonabeau & Meyer, 2001): seed harvester ants (*Messor barbarus*) carrying food back to their nest. Like runners transferring a baton in a relay race, the ants pass food down a chain. The one difference, however, is that the transfer points for the ants are not fixed. That is, an ant carries the food down the chain until it reaches the next ant, and after transferring the food to its nestmate the ant returns back up the chain to receive its next load. The only fixed locations in this operation are the start (the food source) and the end (the nest). This simple approach, dubbed *bucket brigade*, can dramatically increase the efficiency of operations in which work is passed from one person to another.

John Bartholdi of Georgia Tech and Donald Eisenstein of the University of Chicago have studied order pickers at a large distribution center of a major chain retailer. Previously, the warehouse had used a zone approach in which each worker had an assigned area from which to pull products to fulfill an order. One problem with zone approaches is the wide variation in the speeds of different workers—the quickest person could be four times as fast as the slowest. Because of this, zone approaches tend to underutilize the faster people and aggravate the slower ones, who are constantly under pressure to keep up. And, even if every worker were equal, fluctuations in customer demand make it difficult to demarcate the different zones to balance the amount of work. At the distribution center studied, a supervisor had to monitor each aisle to correct for the congestions that inevitably occurred.

Bartholdi and Eisenstein implemented the following rule for each worker: "Continue picking along the flow line until the person downstream of you takes over your work; then head upstream to take over that person's work." The researchers also investigated additional ways in which to maximize productivity. Specifically, should the fastest workers be situated at the start of the flow line, at the end, or somewhere in between? Or should they be located both at the start *and* the finish?

Using ABM, Bartholdi and Eisenstein have proven that the optimum sequence of workers is from slowest to fastest. Using such a scheme, they were able to increase the productivity of the workers at the warehouse studied by more than 30% over the previous zone approach. In addition, the bucket brigade method enables a work line to balance itself; that is, the optimum solution emerges without any intervention by managers. The system is also flexible and robust, easily adapting to any unexpected surges in demand for particular products. Variations of the bucket-brigade approach are being used in the distribution centers of McGraw-Hill; Time Warner Trade Publishing/Little, Brown; Bantam-Doubleday-Dell Distribution; and Blockbuster Music, among others.

Using ABM it was possible to show that this set of DC2 rules generates maximum efficiency. Other examples of applications of ABM to the design of DC2 organizational rules include incentive design (that is, creating an incentive structure that leads to the desired aggregate-level behavior) or the design of risk-mitigating rules to minimize the risk of losses in a corporate or military environment (Bonabeau, 2002a,b).

c. *The ultimate DC2 strategy: swarming and counter-swarming*

In the last two years, swarming has emerged as an important military concept pertinent to DC2 (Bonabeau, et al., 1999; Bonabeau & Meyer, 2001; Arquilla and Ronfeldt, 2000, 2001; Edwards, 2000). Although it has a number of related definitions, in general

> *...swarming occurs when a collection of decentralized, often-diverse units converge on an objective (or a problem) from multiple directions and re-disperse for future action. Swarming suggests the agility to rapidly concentrate the power of a highly networked force in any domain or dimension of warfare to dominate an adversary. Many of the most innovative ideas for U.S. military transformation relate in significant ways to this class of concepts. On the other hand, many threats are much more dangerous when they adopt even crude forms of swarming. As a result, a comprehensive understanding of both the use of, and defense against swarming is critical to effective defense policy* (from the "Swarming: Network Enabled C4ISR" conference, January 2003).

It is generally accepted that swarming concepts have the potential to contribute to U.S. military transformation, that is, the capability to conduct network-centric operations/warfare as envisioned by the DoD Office of Force Transformation and military service transformation plans (see, e.g., Perry et al., 2002).

Indeed, according to Edwards (2000),

> *…the military application of emerging technologies for communications and information processing is likely to change the way military force is managed and applied. One possible change is the reemergence of a doctrine based on swarming, whereby military units organized as networks use dispersed yet integrated operations. This monograph analyzes ten swarming cases throughout history, from Scythian horse archers against a Macedonian phalanx supported by light cavalry, 329-327 B.C., to Somalis surrounding U.S. commandos in a peacekeeping operation in 1993, and applies the conclusions to a discussion of future swarming.*

Or, according to Arquilla and Ronfeldt (2000, 2001), future conflicts will increasingly be clashes of networks rather than hierarchies. According to these authors, swarms are much more than an unorganized mass of individuals – they are a structured and organized way of striking from all directions.

Deleted: ,

> *Swarming is a seemingly amorphous, but deliberately structured, coordinated, strategic way to perform military strikes from all directions. It employs a sustainable pulsing of force and/or fire that is directed from both close-in and stand-off positions. It will work best--perhaps it will only work--if it is designed mainly around the deployment of myriad, small, dispersed, networked maneuver units.*
>
> *This calls for an organizational redesign--involving the creation of platoon-like "pods" joined in company-like "clusters"--that would keep but retool the most basic military unit structures. It is similar to the corporate redesign principle of "flattening," which often removes or redesigns middle layers of management. This has proven successful in the ongoing revolution in business affairs and may prove equally useful in the military realm.*
>
> *From command and control of line units to logistics, profound shifts will have to occur to nurture this new "way of war." This study examines the benefits--and also the costs and risks--of engaging in such serious doctrinal change. The emergence of a military doctrine based on swarming pods and clusters requires that defense policymakers develop new approaches to connectivity and control and achieve a new balance between the two. Far more than traditional approaches to battle, swarming clearly depends upon robust information flows. Securing these flows, therefore, can be seen as a necessary condition for successful swarming* (Arquilla and Ronfeldt, 2000).

It is clear that swarming encompasses a wide range of concepts, ranging from network-centric warfare to fluid warfare tactics. Swarming is relevant to both the design of future forces and the understanding of emerging asymmetric threats that are organized as swarms. Despite the importance of the concept, many military specialists and observers realize how difficult it is to operationalize, first because swarming is the quintessential DC2 concept, and also because it may lead to counter-intuitive emergent phenomena.

ABM provides a powerful answer to this problem. By definition, in ABM each agent in a swarm is represented and modeled, as well as interactions between agents. ABM can therefore capture swarming's emergent phenomena – and can be used to design robust swarming tactics and strategies (see next section). The simple ABM illustration provided in the Emergent Phenomena and Agent-based Modeling section of this paper represents an example of a swarming system. Participants in the game are given simple rules to follow and the aggregate-level result is unpredictable. Using ABM enables modelers to predict the aggregate-level outcome.

Other examples of similar uses of ABM include models of counter-drug modeling by the Bios Group (now NuTech Solutions, Inc.), and by Argonne National Laboratory, sponsored by The Joint Chiefs of Staff, Force Structure, Resources and Assessment Directorate (J-8). Drug traffickers can be viewed as operating as swarms. The ABM not only gives us access to the emergent phenomena that result from the traffickers' swarm behavior, it also enables counter-drug organizations to test DC2, dynamic counter-drugs tactics and strategies.

Another example is the "computer hacker" model developed by Icosystem in collaboration with the US Army Computer Crime Investigation Unit (CCIU). In this model one or more hackers can "swarm" a computer system to perform undesirable actions. Building an ABM of the hacker(s)'s behavior enables the derivation of emergent signatures (that is, in this context, specific log files configurations) which can then be used for intrusion detection and designing DC2 management rules for computer networks. Even in the case of a single or a modest number of hackers, the adverse effects wrought by the hacker(s) can emulate swarming behaviors and thus be analyzed in the same manner as an investigator would look for multiple suspects. It all depends on the paradigm the investigator uses – ABM helps to visualize multiple paradigms and introduce novel lines of inquiry for investigators and analysts (Hunt, 2001).

5. Designing DC2 strategies using ABM

Most of this paper describes how to test DC2 approaches using ABM. However, this assumes that the specifications of the DC2 approach are known. How one might design DC2 specifications becomes a key issue in a context where, by definition, human intuition and experience can be misleading because of emergent phenomena.

Managing and controlling by simple rules is potentially powerful but can also be dangerous, as can be seen in a biological example: army ants. Although army ants are almost blind, they have developed extremely efficient swarming behaviors based on trail laying and trail following: their swarming tactics enables an army ant colony to conduct devastating raids that kill thousands of insects and other arthropods each day. But sometimes a column gets cut off from the rest of the trail because rain has destroyed part of the trail that connected the column to the rest of the raid.

With every individual laying and following, a circular mill may result (see figure 4): a circle of ants continuously following each other round and round in circles until death. Such a mill has been observed in Guyana that measured 1200 feet (365m) in circumference with a circuit time for each ant of about 2½ hours. The mill persisted for two days, with ever increasing numbers of dead bodies littering the route as exhaustion took its toll, but eventually a few workers straggled from the trail thus breaking the cycle, and the raid marched off into the forest. This is clearly an example of undesirable behavior which results from the simple rules individual ants follow. The conclusion is that it is good to know ahead of time that such configurations might occur, and more generally this raises the question of how to design swarming rules and DC2 strategies.

Figure 4
Circular mill in army ants (from Carl Anderson).

Icosystem (Funes and Orme, 2003) has developed an approach to design DC2 rules based on interactive evolution. This approach has been applied to the Aggressor-Defender game described in the Emergent Phenomena and Agent-Based Modeling section, above. The idea is simple: since one may not know ahead of time what to look for, the user is presented a small number of patterns (here, six) generated by different rule sets. In this example, the rule sets are simply characterized by: (1) some people use rule 1 (that is, "try to always have B between A and me") and some people use rule 2 (that is, "try to always be between A and B"); and (2) who interacts with whom (that is, who are A and B for each participant).

When presented with the six patterns resulting from the six different rule sets, the user selects the most "interesting" patterns. An evolutionary algorithm then applies mutation and crossover operators to the rule sets that produced the user's favorite patterns; the offspring rule sets are then used to generate new patterns which are evaluated by the user until desired patterns are observed. As an example, using this approach to evolve interesting patterns in the context of the Aggressor-Defender game, Icosystem found very simple rule sets that produce extraordinary patterns. The user interface used for this work is shown in figure 5.

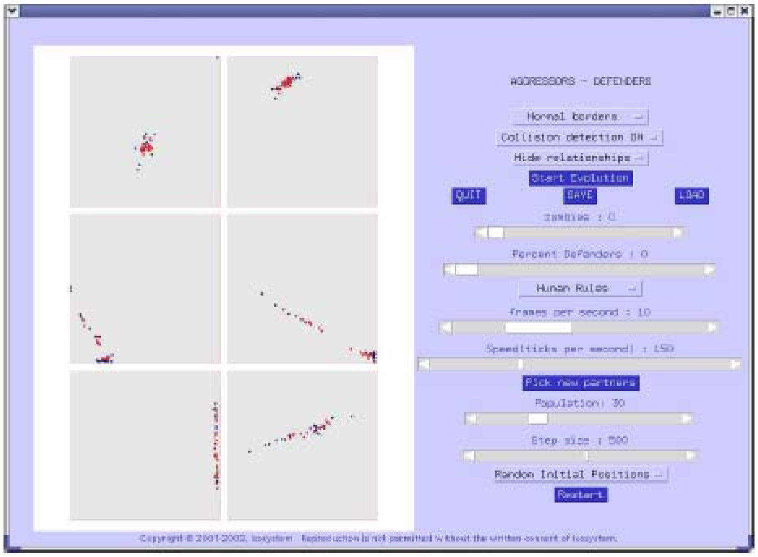

Figure 5
User interface for the interactive evolution system. Clicking on the playground that displays the preferred behavior results in the remaining ones becoming mutants or recombinations of the chosen one (Funes and Orme, 2003).

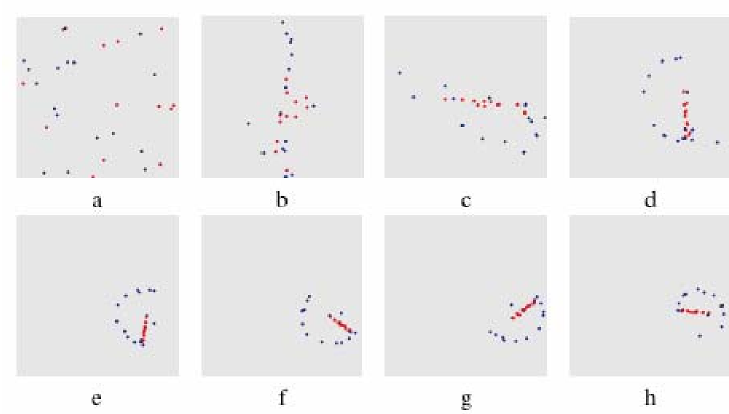

Figure 6
"Chinese streamer." From a random initial placement, a pattern quickly emerges (a-d) and starts turning, stabilizing in a shape with a handle and trailing ribbon which rotates smoothly. The direction of rotation can be clockwise or counterclockwise (as here), presumably depending on the initial positions (Funes and Orme, 2003).

One of the most striking patterns found during this exploration is the "Chinese streamer" presented in figure 6. The reasons why this pattern is striking are twofold: (1) the pattern is complex and highly structured and yet every single agent is only following two simple rules without knowing what the global structure is going to be, and (2) the pattern is robust in the sense that it will form regardless of initial conditions.

After evolving the rule set that produced the Chinese streamer, using ABM combined with interactive evolution, the rules were then tested in the "real world." That is, given real people, the rules, and observation of the emergent pattern, the results were as predicted by the model, as illustrated in figure 7. This very promising result indicates that ABM's general predictive capability can indeed be used to design DC2 management rules.

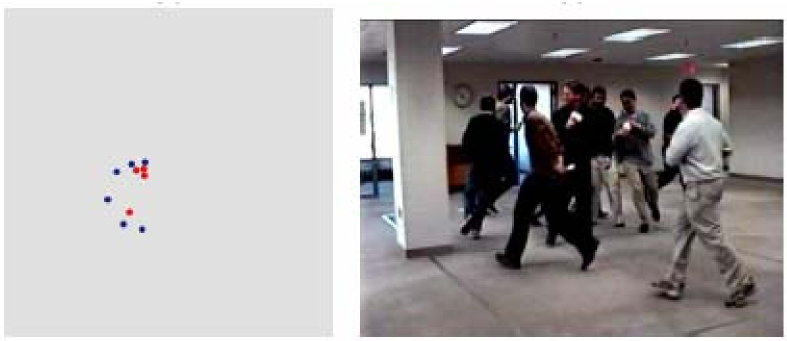

Figure 7
"Chinese streamer" in real life: ABM simulation and "real-world" experiment aimed at testing the rule set (Funes and Orme, 2003).

6. Conclusion

In an increasingly complex environment, confronted with asymmetric threats and fluid, distributed (highly DC2-oriented) enemy organizations, the US military needs to design and test DC2 strategies before it can implement them. That is because although DC2 is an appealing paradigm in the new asymmetric context, it can lead to counter-intuitive and sometimes catastrophic configurations. In this paper we have introduced a new simulation modeling technique, agent-based modeling (ABM) and shown how it can be used to test and even design DC2 strategies.

Agent-based models already exist within the US Defense Department and are undergoing evaluation in a variety of settings. As is the case with any new technology, consistent use and discovery of new ways to use tools that transcend the developers' initial plans often result in even greater innovation than originally conceived. Agent-based models appear to be on the cusp of greater deployability within the US military and may soon be at the heart of significant decision-making processes. Such applications are just in time to be a part of important history being made in the Department and throughout the United States.

References

Arquilla, J. & Ronfeldt, D. (eds) (2001). *Networks and Netwars: The Future of Terror, Crime, and Militancy*. RAND Corporation Publication.

Arquilla, J. & Ronfeldt, D. (2000). *Swarming and the Future of Conflict*. RAND Corporation Publication.

Bonabeau, E., Dorigo, M. & Theraulaz, G. (1999). *Swarm Intelligence*. Oxford University Press.

Bonabeau, E. (2000). Business applications of agent-based simulation, *Adv. Complex Syst*. **3**, 451-461

Bonabeau, E. & Meyer, C. (2001). Swarm Intelligence, *Harvard Business Review*, **May 2001**, 106-114.

Bonabeau, E. (2002a). Agent-based modeling: methods and techniques for simulating human systems, *Proc. Nat. Acad. Sci. USA* **99**, 7280-7287

Bonabeau, E. (2002b). Predicting the Unpredictable, *Harvard Business Review*, **March 2002**, 109-116.

Casti, J. (1997). *Would-Be Worlds: How Simulation is Changing the World of Science* (John Wiley & Sons, New York).

Edwards, S. J. A. (2000). *Swarming on the Battlefield: Past, Present, and Future*. RAND Corporation Publication.

Epstein J. M., Axtell R. L. (1996). *Growing artificial societies: social science from the bottom up* (MIT Press, Cambridge, MA).

Funes, P., Orme, B. (2003). A game of defenders and aggressors: evolving emergent behaviors for simple human agents. Icosystem Tech. Report 0303.

Helbing, D., Farkas, I. & Vicsek, T. (2000). Simulating dynamical features of escape panic, *Nature* **407**, 487-490.

Hunt, Carl W., *Agent Based Evidence Marshaling: Agent-Based Creative Processes for Discovering and Forming Emergent Scenarios and Hypotheses*, Doctoral Dissertation, George Mason University Library, Fairfax, VA, 2001.

Perry, W., Button, R. W., Bracken, J., Sullivan, T. & Mitchell J. (2002). *Measures of Effectiveness for the Information-Age Navy: The Effects of Network-Centric Operations on Combat Outcomes*. RAND Corporation Publication.

Reynolds, C. (1987). Flocks, herds, and schools: a distributed behavioral model. *Computer Graphics* **21**, 25-34.

Still, G. K. (1993). New computer system can predict human behaviour response to building fires, *Fire* **84**, 40-41.

Still, G. K. (2000). Crowd Dynamics, PhD Thesis, Mathematics department, Warwick University.