# Self-organizing Hierarchies in Sensor and Communication Networks

Mikhail Prokopenko*
Peter Wang
Philip Valencia
Ageless Aerospace Vehicle Project
Intelligent Systems
Commonwealth Scientific
and Industrial Research
Organisation (CSIRO)
CSIRO Information and
Communication Technologies
(ICT) Centre
Locked bag 17
North Ryde NSW 1670, Australia
mikhail.prokopenko@csiro.au
peter.wang@csiro.au
philip.valencia@csiro.au

Don Price
Mark Foreman
Anthony Farmer
Ageless Aerospace Vehicle Project
Intelligent Systems
Commonwealth Scientific
and Industrial Research
Organisation (CSIRO)
CSIRO Industrial Physics
Locked bag 17
North Ryde NSW 1670, Australia
don.price@csiro.au
mark.foreman@csiro.au
anthony.farmer@csiro.au

Q1

**Abstract** We consider a hierarchical multicellular sensing and communication network, embedded in an ageless aerospace vehicle, that is expected to detect and react to multiple impacts and damage over a wide range of impact energies. In particular, we investigate self-organization of *impact boundaries* enclosing critically damaged areas, and *impact networks* connecting remote cells that have detected noncritical impacts. Each level of the hierarchy is shown to have distinct higher-order emergent properties, desirable in self-monitoring and self-repairing vehicles. In addition, cells and communication messages are shown to need memory (hysteresis) in order to retain desirable emergent behavior within and between various hierarchical levels. Spatiotemporal robustness of self-organizing hierarchies is quantitatively measured with graph-theoretic and information-theoretic techniques, such as the Shannon entropy. This allows us to clearly identify phase transitions separating chaotic dynamics from ordered and robust patterns.

**Keywords**
Self-organization, sensor networks, impact boundaries, stability metrics, phase transitions, ant colony optimization

## 1 Introduction

NASA's goal of robust or ageless aerospace vehicles requires the development of vehicles that are capable of structural self-assessment and self-repair. The results presented in this article were obtained as part of the joint CSIRO-NASA Ageless Aerospace Vehicle (AAV) research project.

* Corresponding author.

© 2005 Massachusetts Institute of Technology

Previous work in this project developed and examined concepts for integrated smart sensing and communication networks [1, 2, 14, 25, 19]. These concepts are being implemented in a *concept demonstrator*: a hybrid hardware and software system capable of demonstrating principles of an intelligent vehicle health monitoring (VHM) system. Our target modular (multicellular) sensing and communication network is expected to detect, report on and react to impact location, intensity, and damage over a wide range of impact energies, ranging from microparticles to meteoroids. In the face of damage to its parts, the network is expected to self-organise into dynamic and robust hierarchies, so that monitoring and reporting continues as long as possible. To achieve this, we modeled and simulated a multiagent system, made of *cells*, that will not only form a physical shell for an aerospace vehicle, but will also have sensors, logic, and communications (Figure 1). In particular, each cell contains a small number of passive piezopolymer sensors bonded to an aluminum skin panel, and two processors, one of which acquires data from the sensors while the other runs the agent software and controls the communications with its neighboring cells. The processors can be thought of as constituting the cell nucleus. In general, cells may very in shape and size, depending on the surface shape and complexity, and may contain different sensors or multiple sensor types, according to cell location and local requirements. In addition, self-monitoring functions can be divided between those carried out by distributed sensors, intelligent processing, and communication on the skin or within the structure, and those that could be more effectively provided by autonomous mobile nondestructive evaluation (NDE) agents deployed to monitor damage or integrity of the vehicle structure whenever and wherever required.

An aerospace vehicle is likely to be subject to occasional very high-energy impacts. Such impacts are unlikely to be completely avoidable, so that all sensing and processing systems must be redundant. In spite of this, the vehicle will require a high degree of internal coordination. Physically hierarchical systems provide the required coordination, but at the risk of having critical points where damage to a small part of the system may result in a major failure. Such points might be command processors, communications links, or power supply trunks. In general, an integrated VHM system requires multiple functional hierarchies: sensor, communication, inspection, status reporting, repair, and so on. In some cases, the hierarchies will share the same local *hierarchs*, and in others, independent hierarchies will need to evolve. In order to achieve robustness, the system will be constructed so that all or most components and agents have the resources to take any position in any hierarchy, so that if any hierarch fails, its role can be taken by another agent. This leads to a desire for a distributed architecture for the system, where the information is spread over many subsystems, communication and power can be easily rerouted around the damaged area, and loss of any individual piece results in graceful, rather than catastrophic, degradation.

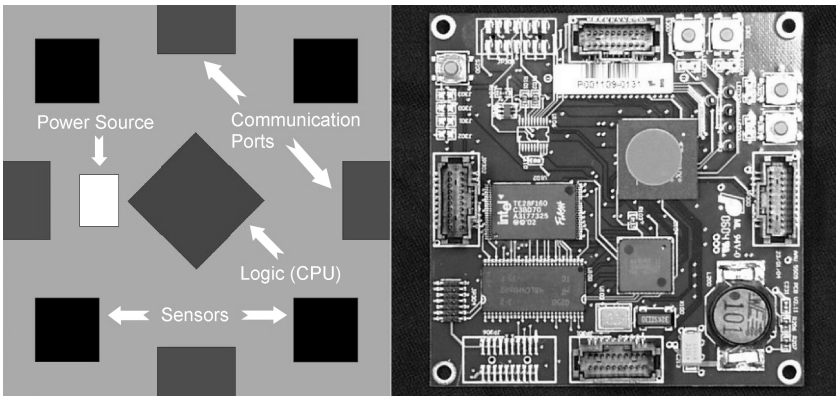


Figure 1. A single square cell with four communication ports. Left: simulator view. Right: a view of a printed circuit board for the agent software sub-module (the connection to the data acquisition submodule is on the other side).

A valuable feature of the distributed approach is the ability of the system to deal with both multiple simultaneous and spatiotemporally distributed real-time events. Hence, the principal task analyzed in this work can be summarized as follows:

> self-organization of stable but reconfigurable information-processing and task-oriented hierarchical structures connecting AAV cells and adapting to environmental changes in real time.

Informally, the emergent hierarchies are expected to deal with *problems that change concurrently with the problem-solving processes.*

## 2 Background and Motivation

In general, the process of self-organization is not random but involves some local information transfer, as pointed out, for example, by Polani [8]. The primary principle that is followed in our work is the *emergence of a global response, including hierarchical organization* as a result of interactions involving transfer of information embedded locally. In other words, without centralized controllers, agents (cells) are expected to self-organize and survive on the basis of local, rather than global, information. Computational modeling at the cellular level in the context of space exploration is not new—a few biological analogies have been investigated in the literature [16]. A concrete and detailed implementation of fault-tolerant circuit synthesis on a self-configurable hardware platform is also provided by the *cell matrix* approach [15, 8]. The approach employs local nodes (*supercells*), performing fault detection, fault isolation, configuration of new supercells, determination of intercell wiring paths, and implementation of the final desired target circuit.

Localized algorithms, in which simple local node behaviors achieve a desired global objective while communicating only with nodes within some neighborhood, have also been investigated in the context of sensor communication networks [10] and amorphous computing [17]. For instance, a simple communication model for describing localized algorithms—*directed diffusion*, proposed by Estrin et al.—includes network nodes that propagate *interests* (sets of labeled data attributes) and establish gradients directing the diffusion of data. As it propagates, data may be locally cached and/or transformed at each node. To a large degree, this idea is inspired by biochemical networks where gradients of different chemical cues play a considerable role in the navigation process, locally embedding positional information. For example, a self-assembly of networklike structures connecting a set of nodes without using preexisting positional information or long-range attraction of the nodes is described by Schweitzer and Tilch [22]—using Brownian agents that are capable of producing different local (chemical) information and responding to it in a nonlinear manner. These agents solve two tasks *in parallel*: (i) the detection of the appropriate nodes, and (ii) the establishment of stable links between them.

Despite some progress, there is a lack of a unifying methodology underlying design of localized algorithms. Our approach to self-organizing hierarchical multi-agent networks is based on an iterative process including the following goals: (a) forward simulation, using localized algorithms and leading to self-organization of hierarchies, and (b) quantitative measurement of spatiotemporal stability of the emerging structures and hierarchies. The information-theoretic metrics obtained at the second step should identify phase transitions in the underlying dynamics and may contribute to evolutionary modeling of the hierarchies [26].

In the context of AAV, we have investigated a number of localized algorithms: formation of an impact boundary with desired characteristics [14, 11], and formation of a spanning tree connecting cells that have detected noncritical impacts [25, 2]. Importantly, these emergent hierarchical structures should be sufficiently stable to perform their tasks while reacting and adapting to change (e.g., new, damaged, or repaired cells). In this article, we describe these multicellular structures in terms of sensor and communication networks—as they emerge on various hierarchical levels. Each "horizontal" level of the main sensor/communication hierarchy may, in general, contain self-organising *internal structural*

*hierarchies* layered on the "plane" orthogonal to the main "vertical axis." For example, each cell may be seen as a self-organized internal hierarchy with varying layers of processing and communication abilities. At the same time, the cell's behavior may contribute to emergence of a multicellular structure—forming at another level of the main dynamical hierarchy. For convenience, we shall differentiate between *layers* of internal hierarchies and *levels* of the main hierarchy. For instance, a well-known dynamical hierarchy exemplified in the "*Ansatz* for dynamical hierarchies" [20] includes water and molecules at level 1, polymers and water at level 2, and micelles (polymer aggregates) and water at level 3. Internally, monomers, polymers, and micelles may also be composed of hierarchical/structural layers—thus, a micelle may have a center of mass and a layered boundary.

## 3 Cellular Structure

> The sense impressions of one-celled animals are not edited for the brain. This is philosophically interesting in a rather mournful way, since it means that only the simplest animals perceive the universe as it is. —Annie Dillard [29]

The first level and the fundamental unit of our main hierarchy is the *cell*. Cells may have varying reactive behaviors embedded in their logic unit. Importantly, the behavior of each cell should be as simple as possible, in terms of the internal logic and the communication policies. *Economy of information* is directly related to the ability to manufacture new and repair or replace damaged cells: the simpler the cell, the easier it is to repair or replace it. A tropistic behavior simply maps sensory inputs to actions, while a hysteretic behavior adds inputs from internal memory. The hysteresis provided by memory inputs addresses some lagging of an effect behind its cause, providing a (temporary) resistance to change that has occurred. For instance, in order to detect a persistent failure in communication with an adjacent cell, the agent needs to filter away spurious occasional miscommunications. Thus, the hysteretic behavior is slightly more deliberate than the tropistic one (exemplified by a simple detection of a missed acknowledgement)—it better situates the agent in time (not only in space) and allows it to better respond to changes. We have investigated a simple hierarchy of the behavioral layers within a cell (tropistic and hysteretic behaviors), while keeping open the possibility of adding or evolving other layers of task-oriented behaviors.

Typically, the damage on the AAV skin caused by a high-energy impact is most severe at the point of impact (an *epicenter*). It will be assumed that not only are the cells at the epicenter destroyed, but the communication capability of the neighboring cells may be reduced. The proximity dependence underlying the probability distribution can be modeled in many ways, and we investigated a range of functions. In general, communication damage propagates out with an exponential decay to a certain radius, and can be characterized by a probability $P_i$ of an error corrupting a message bit $i$, dependent on proximity of the affected communication port to the epicenter:

$$P_i = \frac{1}{2}\left(1 - \frac{d}{\mathrm{R}}\right)^{\epsilon}, \tag{1}$$

where $d$ is the distance between the involved communication port and the epicenter of the impact with the radius $R$, and $\epsilon$ is the decay exponent of the communication loss. Multiple impacts result in overlapping damaged regions. The probability that the whole message containing $n$ bits is corrupted depends on the employed error correction code, and can be estimated as follows:

$$P_{\text{failure}} = \begin{cases} 0 & \text{if} \quad P_i^n \leq r, \\ 1 - P_i^n & \text{if} \quad P_i^n > r, \end{cases} \tag{2}$$

where $r$ is the error correction rate.

A two-dimensional array of cells is modeled in our hybrid concept demonstrator, some cells existing in dedicated hardware (two microprocessors per cell) and some residing within interconnected personal computers (a number of cells per PC) [19, 2]. We also used a standalone AAV simulator capable of simulating some simple environmental effects such as particle impacts of various energies [14]. In the AAV simulator, cells are represented as objects (polygons) on a two-dimensional plane (e.g., Figure 2), where they asynchronously interact *only* with their immediate neighbors in a von Neumann neighborhood, through connected (geometrically overlapping) communication ports.

## 4 Multicellular Spatial Hierarchies—Impact Boundaries

In this section, we describe emergent hierarchical multi-agent structures capable of tracing *impact boundaries* in the presence of cell failures and connectivity disruptions. On the one hand, it is desirable that an impact boundary, enclosing damaged areas, form a continuously connected closed circuit. This circuit may serve as a reliable communication pathway around the impact-surrounding region within which communications are compromised. Every cell on a continuously connected closed circuit must always have two and only two neighbor cells, designated as the circuit members (circuit neighbors of this cell). A non-closed continuously connected circuit (a chain) may have two terminal cells, each with only one circuit neighbor. On the other hand, a continuously connected closed impact boundary provides a template for repair of the impact-surrounding region, uniquely describing its shape (Figure 3). Both these functionalities of impact boundaries can be contrasted with noncontinuous *guard walls*, investigated by Durbeck and Macias [8], that simply isolate faulty regions of the cell matrix, without connecting elements of a guard wall in a circuit. Importantly, impact boundaries should be robust to communication failures caused by proximity to the impact, due to high probabilities $P_{\text{failure}}$, and without knowing distances between each cell and the epicenter. In short, the aim is to achieve spatiotemporal stability of impact boundaries.

In the remainder of this section we shall describe a series of emergent structures, progressing towards stable impact boundaries and impact-surrounding regions. We shall begin in Section 4.1 with policies and behaviors producing a pair of internal layers: a *scaffolding* region containing the cells that suffered significant communication damage, and a *frame* boundary supported by the scaffolding region and separating it from the cells that are able to communicate to their normal functional capacity. The frame boundary provides neither a closed nor a continuously connected circuit, but rather forms a regular stable frame for the desired *closed* boundary. Section 4.2 describes the policies needed for self-organisation of such a closed impact boundary, connecting the cells on the frame boundary into a continuous closed circuit by identifying their circuit neighbors. These internal layers (scaffolding, frame, and closed boundary) completely define an impact-surrounding region as a layered spatial hierarchy. Finally, Section 4.3 "zooms" inside the impact-surrounding region and identifies its emergent properties, including a recovery membrane stabilizing this hierarchy.

### 4.1 Frame Boundary

Our first target is the formation of a preliminary frame boundary. In order to support a closed, continuously connected circuit, a *regular* frame should not be too thin (a scaffolding cell must not be adjacent to a normal cell) and should not be too thick (there must be no frame cells in the direction orthogonal to a local frame fragment). To achieve this, we employed one of the most basic communication policies:

(i) every cycle each cell sends an "OK" message to all its neighbors;

together with the most basic tropistic behaviors embedded in each cell's logic unit:

(ii) if all neighbor cells failed to communicate the "OK" message this cycle, switch to the scaffolding state $S_s$;

(iii) if at least one neighbor cell failed to communicate the "OK" message this cycle and at least one neighbor cell did communicate the "OK" message this cycle, switch to the frame boundary state $S_f$.

To reiterate, the internal scaffolding region is intended to contain all the cells that cannot communicate reliably. The frame, on the other hand, is an inner layer of normal cells that are able to communicate reliably among themselves. Figure 2 illustrates the situation: white cells are destroyed, dark-gray cells indicate scaffolding state, and black cells self-organize in the frame boundary; normal cells (neither scaffolding nor boundary) are shown in background color.

However, the combination (i)–(iii) may still produce some irregularities in the frame. Figure 2 (left) shows some of the irregularities: (a) on the lower side, a scaffolding cell $\alpha$ is next to a normal cell; (b) there are protruding frame cells on all other sides, emerging in directions orthogonal to their respective local frame fragments. These abnormalities occur because communication failures are *asymmetric*: in general, if a message from one cell reaches another cell, the reverse may not be the case. For example, the scaffolding cell $\alpha$ [Figure 2 (left)] did not receive the "OK" message from any of its neighbors. However, the normal cell directly below $\alpha$ received the "OK" message from all its neighbors including $\alpha$, while not being able to transmit the "OK" message to the cell $\alpha$. A similar asymmetry explains thickening of the frame. To resolve irregularities illustrated by these asymmetries, we augmented the communication policy by acknowledgment messages:

(iv) upon receiving the "OK" message each cell replies with an "Acknowledgment" message,

and replaced behavior (ii) with the following:

(v) if all neighbor cells failed to communicate the "Acknowledgment" message this cycle, switch to the scaffolding state $S_s$, and stop sending messages.

The policy (i), (iii)–(v) allows each cell to go into "silent mode," in which it may receive but chooses not to transmit if it detects that its own communication channel has been compromised. This results in an emergent regular frame boundary surrounding the scaffolding region [Figure 2 (right)].

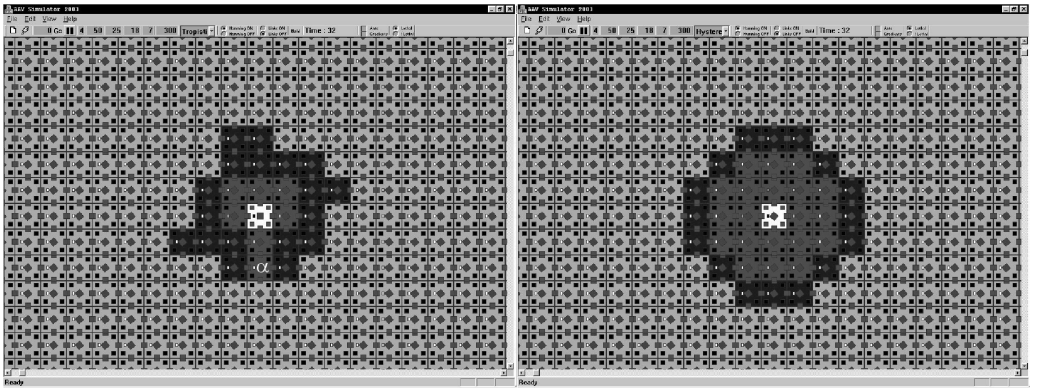


Figure 2. A white cell is destroyed; dark-gray cells form *scaffolding*; black cells form the *frame*. Left: An irregular frame boundary. Right: A regular frame boundary. There are no scaffolding (dark-gray) cells neighboring a normal (gray) cell.

Nevertheless, on some occasions the boundary is erratic and unstable—some cells frequently change their states between frame, scaffolding, and normal. The reason for this spatial and temporal instability is that a simple detection of a missed acknowledgement may sometimes be insufficient—especially when the involved cells are on a periphery of the damaged area. In other words, in order to detect a persistent failure in communications with an adjacent cell, the cell needs to filter away spurious occasional miscommunications. This is not possible with tropistic algorithms. However, behaviors (iii) and (v) can be implemented as hysteretic: the failure or success of a communication with each neighbor cell can be judged not instantaneously ("this cycle") but over a period. More precisely, a failure or a success is detected if the percentage of lost messages during the last $\mu$ cycles is greater than a certain threshold. The length $\mu$ limits the cell-to-cell communication history kept by each cell. The hysteretic version of the algorithm provides a degree of stability in the treatment of communication connections between any two cells, and results in a stable and regular self-organizing frame boundary around the damaged area.

### 4.2 Closed Boundary

The scaffolding region and frame boundary are the first two hierarchical or structural layers leading to the desired continuously connected boundaries. In order to generate links between disconnected cells on the frame boundary, we needed to introduce the following communication rules and behaviors:

(vi) if the cell state is $S_f$, and there are at least two communicating neighbors $\beta_1$ and $\beta_2$, switch to the closed state $S_c$;

(vii) if the cell state is $S_c$, then (1) determine a (the) cell $\alpha$ that failed to communicate; (2) determine two communicating neighbor cells $\beta_1$ and $\beta_2$ nearest to the cell $\alpha$ on opposite sides (clockwise and counterclockwise relative to $\alpha$); (3) establish a boundary link between ports to $\beta_1$ and $\beta_2$ by designating $\beta_1$ and $\beta_2$ as circuit neighbors of the cell; (4) map the directions to $\alpha$, $\beta_1$ to a direction $\gamma_1$, and the directions to $\alpha$, $\beta_2$ to a direction $\gamma_2$; and (5) send a "Connect($\tau$, $\gamma_i$)" message to each circuit neighbor $\beta_i$, $i \in \{1, 2\}$, with a "*time-to-live*" parameter $\tau$;

(viii) upon receiving a "Connect($\tau$, $\gamma$)" message from the cell $\delta$, if the cell state is not $S_c$, then (1) switch to the state $S_c$; (2) determine the communicating neighbor cell $\beta$ nearest to the cell $\delta$ in the direction $\gamma$; (3) establish a boundary link between ports to $\delta$ and $\beta$ by designating $\delta$ and $\beta$ as circuit neighbors of the cell; (4) if $\tau > 0$, map the directions to $\delta$ and $\beta$ to a direction $\gamma$; and (5) send a "Connect($\tau - 1$, $\gamma$)" message to the circuit neighbor $\beta$.

Let us exemplify these rules with interactions producing a connected boundary fragment $\beta_1$-$\delta$-$\beta_2$-$\beta$. Rule (vi) enables the process of establishing boundary links between ports that would connect circuit neighbors in the possible directions (top, down, left, right). For example, a square cell $\delta$ that has reliable communications with its left and right neighbors may establish a left-right link. This link becomes part of the boundary if, for example, the top neighbor (e.g., cell $\alpha$) fails to communicate. This is detected by rule (vii), which recognizes potential circuit neighbors on the left (say, $\beta_1$) and on the right ($\beta_2$) as closest to $\alpha$ on opposite sides. Each pair of directions (e.g., to $\alpha$ and $\beta_2$) is uniquely mapped to a direction (say, $\gamma_2$), and this map is encoded in each cell's logic. The connection process is triggered in the respective direction. Rule (viii) simply propagates (or completes) the process in the communicated direction: for instance, when cell $\beta_2$ receives from $\delta$ the message directed to the latter's right direction, it determines the other circuit neighbor that is as close to this direction as possible (say, $\beta$). The time-to-live parameter $\tau \geq 1$ determines the maximum amount of time the message may live in the network, effectively setting a limit on the number of cells the message may pass through before being discarded.

The policy (i),(iii)–(viii) achieves the desired robustness and continuity of self-organizing impact boundaries (layer 3) for a variety of cell shapes (triangular and square), impact energy dissipation

profiles, and communication damage probability models in Equations 1 and 2, that is, varying the exponential decay $\epsilon \in [1, 7]$ and the error correction rate $r \in [0, 1/8]$. The hysteretic behavior handles well the more demanding cases of exponential decreases $\epsilon > 1$ [Figure 3 (left)]. The length $\mu$ of the hysteretic filter (i.e., the cell-to-cell communication history) is an important parameter, directly affecting spatial and temporal stability of self-organizing impact boundaries. Our experiments, described in the next section, verified that larger values of $\mu$ lead to more stable boundaries, while $\mu = 0$ (no hysteresis at all—purely tropistic behavior) could not produce stable boundaries [Figure 3 (right)].

### 4.3 Emergent Membranes and Recovery Mode

> Men, it has been well said, think in herds; it will be seen that they go mad in herds, while they only recover their senses slowly, and one by one. —Charles Mackay [30]

The previous section described multicellular structures—impact boundaries—emerging at a higher level of the main multi-agent hierarchy with distinct higher-order properties of being *closed* and *continuous*. It could be argued that frame and closed boundaries emerge at level 2 (analogous to polymers) above the cell level, and the whole *multilayered impact-surrounding region* emerges at level 3 (analogous to micelles). In other words, formations that emerged at different levels of the main sensor-communication hierarchy may also serve as structural layers within a particular object at a given level. As an emergent structure, the impact-surrounding region has unique third-order properties, such as having an *inside* and an *outside*, and importantly, a *recovery* property.

To support these conjectures, let us consider the condition (v) introduced in order to break asymmetry between neighboring cells: whenever a cell is able to transmit data but not receive, or vice versa, it decides not to communicate at all. As mentioned, this behavior is effective in achieving stable impact boundaries. However, the cells that have stopped transmitting messages may need to resume communications under certain conditions—for example, when a repair action is initiated and their neighbors, in fact, are ready to receive communications (i.e., when the cause of asymmetry is eliminated). The conditions for resumption of communications have to be precise so that they are not reacted upon prematurely, interfering with boundary formation. A variant of these *recovery* conditions is given below:

(ix) if the cell state is $S_s$, and there are no miscommunicating neighbors (during period $\pi$), switch to the recovery state $S_r$ and resume sending messages.

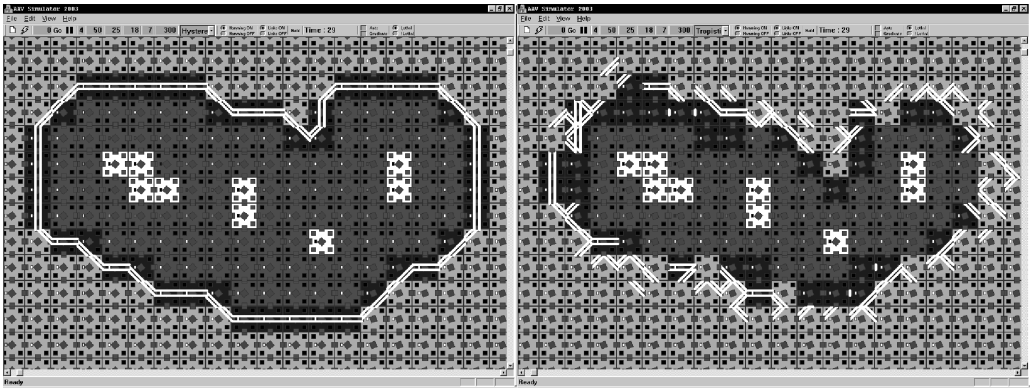

Figure 3. White cells are destroyed; dark-gray cells form scaffolding; black cells form the frame. Boundary links are shown as white double lines. Left: A stable impact boundary ($\mu = 10$). Right: An unstable impact boundary ($\mu = 0$).

(x) if the cell state is $S_s$, and there are no communicating neighbors (during period $\pi$), switch to the recovery state $S_r$ and resume sending messages.

The combination of the shutdown condition (v) and the recovery conditions (ix) and (x) makes the problem of designing stable impact boundaries particularly difficult. On the one hand, condition (v) disables asymmetric messages, making the task of boundary formation easier. On the other hand, conditions (ix) and (x) may enable asymmetric messages, potentially destabilizing the boundary.

The solution to this problem is provided by emergent spatiotemporal structures—recovery membranes—that separate the boundaries from recovering cells. A recovery membrane belongs to level 2, and always forms inside the closed boundary and outside the recovering area. Importantly, membrane forms within the impact-surrounding region, proving the point that the latter has an inside and an outside. Interestingly, unlike scaffolding and frame boundary, the membrane is not a designated state into which a cell can switch. In other words, membrane cells shut down their communications, following behavior (v) like other scaffolding cells, but do not resume communications, because behaviors (ix) and (x) are not applicable, as there are some (but not all) miscommunicating neighbors (determined hysteretically, over the recovery-delay period $\pi \geq 1$). Without a membrane, the cells on the frame boundary would be confused by intermittent messages from scaffolding cells attempting recovery. Figure [4 (left)] illustrates a recovery membrane shown with dark-gray color, while the recovering cells are shown in darker shade of white. It is quite obvious that the membrane cells separate the scaffolding cells from the frame boundary.

If a repair is initiated and a healthy cell replaces a scaffolding cell (including one in the membrane), then the recovery mode results in shrinking membranes and impact-surrounding regions. Figure 4 (right) shows a scenario where eight healthy cells were removed from the the top right corner of the entire array and used in replacing eight damaged cells in the top right corner of the impact-surrounding region. The repair action, used in the scenario, was not initiated by the system itself at this stage (the 16 involved cells were manually transported), and was done only to illustrate the recovery mode and potential for self-repair.

## 5 Spatiotemporal Stability and Phase Transitions

It seems that an essential condition for crises is to be found in the existence of a highly developed system of communications and the spreading of a homogenous [*sic*] mentality

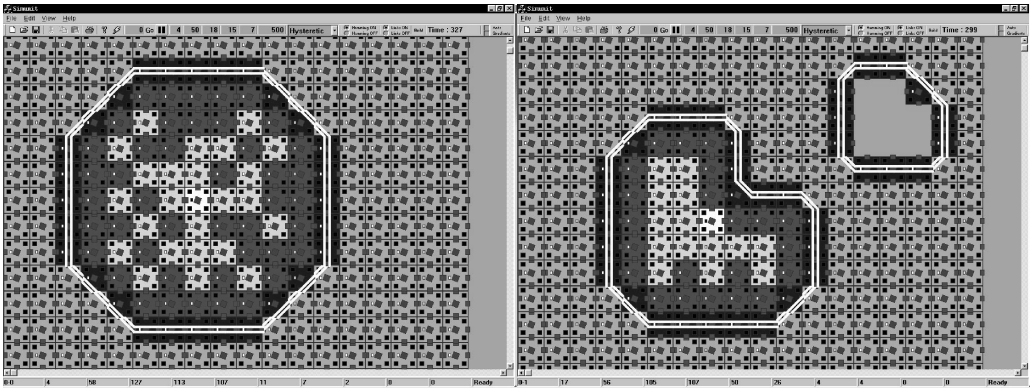

Figure 4. A white cell at the epicenter is destroyed. Left: Scaffolding cells around the impact attempt recovery and are shown in a darker shade of white. A recovery membrane, shown in dark gray, absorbs recovering cells and separates them from the frame shown in black. Right: Another impact-surrounding region is shown in the array's center. Its top right corner is repaired with the cells taken from the entire array's top right corner.

> over vast areas. But when the hour and the right material are at hand, the contagion spreads with the speed of electricity over hundreds of miles, and affects the most diverse populations, which hardly know each other. The message flies through the air and they all suddenly agree on that one issue, if only a sulky admission that "there's got to be a change".—Jakob Burckhardt [31]

The analysis of spatiotemporal stability of self-organizing impact boundaries with respect to the parameter $\mu$ is the primary subject of our investigation in this section, following the results reported in [11]. In particular, we characterize the dynamics of multicellular impact boundaries in terms of generic graph-theoretic and information-theoretic properties, and pinpoint certain phase transitions separating chaotic dynamics from ordered and robust patterns.

One information-theoretic measure targeting the complexity of a multi-agent system was investigated by Wuensche [28], who characterized rule spaces of one-dimensional cellular automata with the frequency distribution of the Shannon entropy of rules. The input entropy settles to fairly low values for ordered dynamics, but fluctuates irregularly within a narrow high band for chaotic dynamics. For the complex CA, the input entropy generally settles onto a short attractor cycle, where order and chaos may predominate at different times, causing the entropy to vary. A measure of the variability of the input-entropy curve is its variance or standard deviation, calculated over time. Wuensche has convincingly demonstrated that only complex dynamics exhibits high variance of input entropy, leading to automatic classification of the rule space. Importantly, the peak of the input-entropy variance points to a phase transition as well, indicating the edge of chaos (complexity). Langton [13] also observed that the individual cell entropy showed a discrete jump from low to high. This evidence pointed to a first-order phase transition, similar to that observed between the solid and fluid phases of matter. Interestingly, the average mutual information, defined as a function of individual cell entropies for a particular value of the parameter $\lambda$ (the fraction of rules in which the new state of the cell has a given property), has a distinct peak at this point, indicating a phase transition from order to chaos.

We would like to point out, however, that the input entropy of the CA rule space relies purely on local state variables that can be associated with each cell, and does not directly capture *inter-agent* connections and interactions. In other words, it traces diversity of rules used over time and identifies temporally persistent configurations, rather than spatial connectivity among neighboring cells. More precisely, if one needs to capture emergence and stability of a specific spatial pattern (e.g., a continuously connected boundary, a maximally connected subgraph, or a desired geometric or topological shape), then another type of measure is needed.

It is well known that one feasible measure of a network's heterogeneity is given by the entropy defined through the link distribution:

$$H(P) = \sum_{k=1}^{m} P_k \ \log(1/P_k). \tag{3}$$

The latter can be defined via the simple degree distribution—the probability $P_k$ of having a node with $k$ links. Ultimately, the amount of correlation between nodes in the graph can be calculated via the mutual information measure (the information transfer) as

$$I(P) = \sum_{k=1}^{m} \sum_{k'=1}^{m} P_{k,k'} \ \log \frac{P_{k,k'}}{P_k P_{k'}}, \tag{4}$$

where $P_{k,k'}$ is the joint probability of connected pairs [23].

Another indicator of phase transitions in the graph connectivity is the size of the largest maximally connected subgraph (LCS) and its variance or standard deviation obtained across an ensemble of graphs. Critical changes in the connectivity of a directed graph occur as the number of edges increases—the size of the LCS rapidly increases as well and fills most of the graph, while the variance in the size of the LCS reaches a maximum at some critical point before decreasing [9, 12]. This variance approximates the mutual information contained in the system, combining both individual cells' diversity and their interdependence (correlation). In general, the mutual information defined by Equation 4 is a better, more generic measure of dependence: correlation functions, underlying the variance in the size of the LCS, "measure linear relations whereas mutual information measures the general dependence and is thus a less biased statistic" [23].

However, for our purposes it is more practical and sufficient to characterize the diversity of and interdependence within impact boundaries (corresponding to different values of the parameter $\mu$) through the variance in the size of boundary fragments over time. Typically, when a critical parameter changes, individual diversity and interdependence move in opposite directions and clash at some transition point—at the *edge of chaos*, where the system dynamics exhibits the most complexity. One may expect that increasing the length $\mu$ of the cell-to-cell communication history in our experiments would also increase correlation and interdependence between neighbor cells, and decrease their diversity. Therefore, a phase transition in the parameter space may be expected at some critical value of $\mu$, where the variance in the size of boundary fragments peaks. Thus, a metric based on this variance may capture *spatial* aspects of desired emergent patterns (connectivity), while a metric based on the entropy of local agent variables (rules, coordinates, velocities, etc.) may capture *temporal* persistence of emergent behavior.

### 5.1 Spatial Stability

The analysis of the spatial stability of impact boundaries is based on the concept of a connected boundary fragment (CBF). A CBF is simply a set $F$ of cells in the closed state $S_c$ such that every cell in $F$ is connected with at least one other cell in $F$, and there exists no cell outside $F$ that is connected to at least one cell in $F$ (an analogue of a maximally connected subgraph or a graph component). We carried out 63 experiments (3 for each value of $\mu$ from 0 to 20). Each experiment simulates an impact at a predefined cell, and lasts 500 cycles; the first $2\mu$ cycles are excluded in order to avoid penalizing longer history lengths. At each time point (cycle) we calculate the average size $H_{sp}(\mu)$ of CBFs in self-organizing impact boundaries, and its standard deviation $\sigma_{sp}(\mu)$ over time (500 cycles). Figure 5 plots data of five such experiments: $H_{sp}(\mu)$ for $\mu = 0$ (left), $\mu = 4$ (strongly fluctuating), and $\mu = 8$, $\mu = 12$, $\mu = 16$, $\mu = 20$ (right). The impact boundary in this case has the size 24, and the hysteretic behavior with lengths $\mu > 4$ leads to more stable self-organization—the four rightmost plots are almost straight lines. On the other hand, the cells without any hysteresis ($\mu = 0$) are not capable of self-organization at all—an average CBF contains about five cells. The intermediate case ($\mu = 4$) is most unstable, however! Sometimes it creates boundaries twice as long as needed, and sometimes it collapses into a lot of small fragments. In other words, the hysteretic behavior leads to order, while tropistic behavior, using a direct mapping from communication inputs to actions ($\mu = 0$), produces chaotic dynamics. In the region $\mu \in [1, 4]$ we observe *complex* dynamics, and the plot of the mean standard deviation $\sigma_{sp}(\mu)$ of average size $H_{sp}(\mu)$, shown in the foreground Figure 5, makes the phase transition apparent—between $\mu = 4$ and $\mu = 5$.

It is important to realize that the phase transition is observed in the parameter space of the communication history $\mu$, and not within a given graph or network. In other words, for a specific length $\mu$, boundary fragments do not gradually grow until they form a whole. They either rapidly link together during the first few cycles (approximately $2\mu$) in the ordered cases with the longer lengths $\mu$, or do not form a boundary at all in the chaotic cases with the shorter lengths $\mu$. At the edge of chaos (in terms of $\mu$), boundaries exhibit both patterns, never settling to any specific dominant regime. Thus, we would like to point out a subtle difference between (a) the variance in the size of the average CBF, calculated over time for a given $\mu$, and (b) the variance in the size of the LCS, calculated over multiple random graphs for a given degree (link) probability. Nevertheless, our

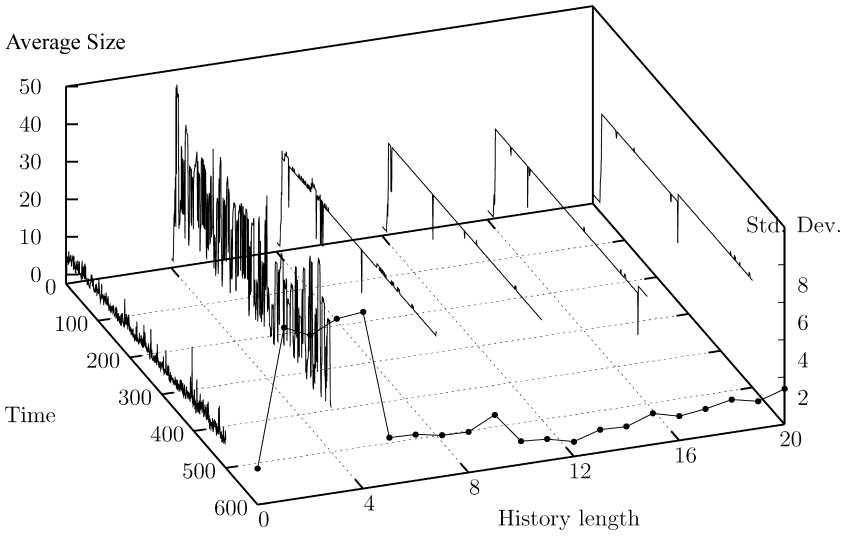


Figure 5. Background (left *z* axis): average size $H_{sp}(\mu)$ of CBFs, for $\mu = 0$ (left), $\mu = 4$ (strongly fluctuating), and $\mu = 8$, $\mu = 12$, $\mu = 16$, $\mu = 20$ (right). Foreground (right *z* axis): standard deviation $\sigma_{sp}(\mu)$ of average size $H_{sp}(\mu)$.

experiments demonstrated that a stability metric based on the variance in the size of the CBF captures *spatial connectivity* in impact boundaries. By itself, however, the metric would be insufficient: a continuous boundary may change its shape over time, without breaking into fragments, and keeping the size of CBF almost constant.

### 5.2 Temporal Stability

In order to analyze temporal stability, we consider state changes in each cell at every time step. There are six symmetric boundary links possible in each square cell, connecting ports left-right, top-bottom, left-top, and so on (Figure 3). Thus, there are $2^6$ possible boundary states (permutations of the boundary links), including "no boundary," and $m = 2^{12}$ transitions. These transitions describe space-time dynamics inside a one-cell neighborhood (containing just the cell itself), and exemplify a more general case of $p$-cell neighborhoods (the von Neumann neighborhood with $p = 5$, the Moore neighborhood with $p = 9$, etc.). For each value of the parameter $\mu$ (varying from 0 to 20), we calculate the entropy $H_{\text{temp}}(\mu)$ of a particular frequency distribution $X_i^t(\mu)$, where $t$ is a time step, and $i$ is a cell transition index: $1 \leq i \leq m$. Analogously to the analysis conducted by Wuensche [28], we define the input entropy as

$$H_{\text{temp}}(\mu) = -\sum_{i=1}^{m} \frac{X_i^t(\mu)}{n} \log \frac{X_i^t(\mu)}{n},$$

where $n$ is the system size (the total number of cells), and $X_i^t(\mu)$ is the number of times the transition $i$ was used at time $t$ across the system. The temporal entropy $H_{\text{temp}}(\mu)$ can also be interpreted as the joint state transition entropy $H(S^t(\mu), S^{t+1}(\mu))$, where $S^t$ is the state of the cell at time $t$. Figure 6 plots the temporal entropy $H_{\text{temp}}(\mu)$ for different history lengths $\mu$. The plots for all $\mu > 4$ cases are almost straight lines—there is practically no variation in the entropy values, as the boundary is stable. We conducted three experiments for every value of $\mu$, each time calculating the standard deviation $\sigma_{\text{temp}}(\mu)$ over 500 cycles. The plot of the mean $\sigma_{\text{temp}}(\mu)$ is shown in the foreground of Figure 6, clearly identifying a phase transition for some critical value of $\mu$ between $\mu = 4$ and $\mu = 5$.

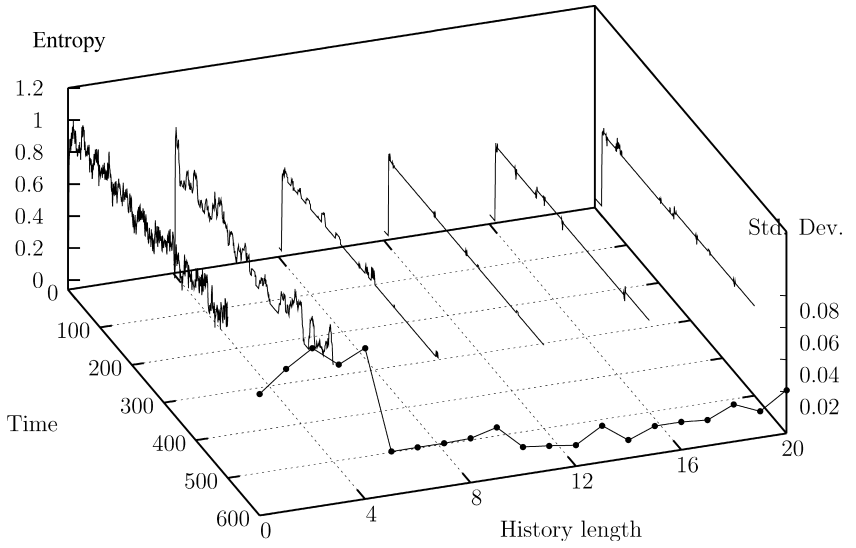


Figure 6. Background (left $z$ axis): temporal entropy $H_{\text{temp}}(\mu)$ for $\mu = 0$ (left), $\mu = 4$ (strongly fluctuating), and $\mu = 8$, $\mu = 12$, $\mu = 16$, $\mu = 20$ (right). Foreground (right $z$ axis): standard deviation $\sigma_{\text{temp}}(\mu)$ of temporal entropy $H_{\text{temp}}(\mu)$.

This variance, in a sense, captures mutual information contained in boundary's dynamics, combining both individual rules' diversity and their interdependence (correlation). In particular, as the length $\mu$ increases, rules in neighbor cells become more correlated and less diverse.

These experiments demonstrated that a metric modeled after the input entropy $H_{\text{temp}}(\mu)$ captures the temporal stability of impact boundaries, but not the spatial stability of continuous and closed circuits. Put simply, $H_{\text{temp}}(\mu)$ may be very low and may vary only slightly for temporally stable patterns that keep their disconnected shape. Together, the temporal and spatial metrics capture both important aspects—steadiness and continuity of impact boundaries.

## 6 Impact Networks

Impact boundaries and regions emerge as a result of local and short-range interactions. In other words, an impact-surrounding region with third-order properties is a spatial aggregation. However, as pointed out by Rasmussen et al. [20], "aggregation is not the only way higher-order structures may arise. . . . [I]n networks, there may be functionally higher-order properties that are topologically distributed." In this section, we briefly describe emergence of another hierarchical level—a robust reconfigurable network connecting remote AAV cells that belong to a specific class [25, 2]. One subtask (decentralized inspection) may require an *impact network* among cells that registers impacts with energies within a certain band (e.g., noncritical impacts). In these scenarios the communication links even in the immediate proximity to the impact point are likely to remain intact. The self-organizing impact network can facilitate fast inspections of impact areas and identification of high densities of impacts (e.g., meteor showers). Another subtask (decentralized repair) may need a self-organizing network connecting local hierarchs responsible for coordination and routing of repair resources. A shortest or *minimum spanning tree* is often required in order to efficiently solve these decentralized subtasks. Importantly, the problem's graph representation changes over time due to the appearance of new nodes (e.g., noncritical impacts), while the problem's properties change due to varying connection costs (e.g., critical impacts destroying existing paths). In short, the problem changes concurrently with the problem-solving process again.

These factors suggested that the problem of forming minimum spanning trees on the AAV skin can be efficiently tackled by ant colony optimization (ACO) algorithms proposed and enhanced over recent years by Dorigo and his colleagues [5, 4, 6, 7]—rather than distributed dynamic programming (Bellman-Ford) algorithms or reinforcement learning techniques such as backtracking. An overview of the ACO meta-heuristic and its applicability can be found in [7]. Essentially, the ACO algorithms use the ability of agents to indirectly interact through changes in their environment (*stigmergy*) by depositing pheromones and forming a pheromone trail. They also employ a form of *autocatalytic* behavior—*allelomimesis*: the probability with which an ant chooses a trail increases with the number of ants that chose the same trail in the past. The process is thus characterized by a positive feedback loop [5, 6]. The most interesting aspect of this autocatalytic process is that finding the shortest path around an obstacle is an emergent property of the interaction between the obstacle shape and the ants' distributed behavior.

Each cell (node) in the network stores an ant-routing table, used in determining which neighbor cell should be chosen by an incoming ant packet to continue their travel. The ant-routing table $a_{i,j}(t)$ of node $i$ with respect to all its neighbor nodes $j$, contained in the neighborhood $N_i$, is usually obtained by the following functional composition of pheromone trails $\varphi_{i,j}(t)$ and local heuristic values $\eta_{i,j}(t)$:

$$a_{i,j}(t) = \frac{[\varphi_{i,j}(t)]^{\alpha} \; [\eta_{i,j}(t)]^{\beta}}{\sum_{l \in N_i} [\varphi_{i,l}(t)]^{\alpha} \; [\eta_{i,l}(t)]^{\beta}} \quad \forall j \in N_i,$$

where $\alpha$ and $\beta$ are two parameters balancing the relative weights of pheromone trail and heuristic value. When solving the traveling-salesman problem (TSP), the local heuristic $\eta_{i,j}(t)$ is the visibility, determined as the inverse distance between nodes $i$ and $j$, and indicating that closer nodes are more strongly preferred. The transition of the ant $k$ from node $i$ to node $j$ at time $t$ is driven by the following transition probabilities:

$$p_{i,j}^{k}(t) = \frac{a_{i,j}(t)}{\sum_{l \in N_i^k} a_{i,l}(t)}, \tag{5}$$

where $N_i^k$ is the feasible neighborhood of node $i$ for ant $k$, determined by the problem constraints such as current obstacles.

The trail intensity $\varphi_{i,j}(t)$ on the path between nodes $i$ and $j$ gives information on how many ants have chosen the path in the past. The trail intensity is updated each time an ant agent $k$ passes through the node:

$$\varphi_{i,j}(t) = \varphi_{i,j}(t) + \Delta\varphi_{i,j}^{k}(t).$$

At the beginning (in some ACO variants, at the end) of each cycle, the pheromone evaporates at the rate $\rho \in [0, 1]$:

$$\varphi_{i,j}(t) = (1 - \rho)\, \varphi_{i,j}(t).$$

For simplicity, we shall use $\varrho = 1 - \rho$ to denote the pheromone retention coefficient: higher $\varrho$ corresponds to higher retention of pheromones.

Ants are usually generated with certain frequency, or at a given interval $T$ (a given number of cycles), and terminate after a fixed number of cycles greater than $T$. The parameters $\alpha$ and $\beta$ determine the relationship between intensity and visibility. For example, setting $\alpha = 0$ results in classical stochastic search algorithms, while the choice $\beta = 0$ ignores the local domain knowledge and leads to a pure autocatalytic process. These two extreme options exhibit convergence to suboptimal solutions.

The algorithm presented in [25] was based on a hybrid method of establishing impact networks, using a single *impact gradient field* (IGF) and a dead-reckoning scheme (DRS), complementing the autocatalytic process of antlike agents. Here we shall describe a main variant of this algorithm, without an IGF, and relying only on DRS. The ants are implemented as communication packets themselves, so the policies are implemented via appropriate message passing, where the cells are responsible for unpacking the packets, interpreting them, and sending updated packets further if necessary. Thus, ants cannot move into the cells with damaged (or shut-down) communication links, so regions enclosed by the self-organising boundaries described in Section 4 form obstacles, and the ants are expected to find the shortest paths around them using positively reinforced pheromone trails. Agents' exploring behavior includes the following:

(E1) each impact node generates a number of exploring ants every $T$ cycles; each ant has a time-to-live counter, decremented every cycle;

(E2) an exploring ant performs a random walk until either (a) another impact node is found, or (b) the ant has returned to the home impact node, or (c) the ant can move to a cell with a nonzero trail intensity;

(E3) if an exploring ant can move to a cell with a nonzero trail intensity, the destination cell is selected according to the transition probabilities (Equation 5);

(E4) at each step from cell $i$ to cell $j$, an exploring ant updates the $x$ and $y$ shift coordinates from the home node (initially set to 0).

The DRS is inspired by Tunisian desert ants studied by Wehner and Srinivasan [27], and assumes that each moving agent (ant) remembers the $x$ and $y$ shift coordinates from the home node. These coordinates are relative; they simply reflect how many cells separate the ant from the home node in terms of $x$ and $y$ at the moment, and should not be confused with a taboo list of an ACO agent containing all visited nodes in terms of some absolute coordinate or identification system. The DRS enables the agents to head home when another impact node is located:

(R1) when another impact node is found, the exploring ant switches to a return state, remembers the ratio $g = y/x$ corresponding to the found node's coordinates relative to the home node, and starts moving back to the home node by moving to cells where the $y$ and/or $x$ shift coordinate(s) would be smaller and their ratio would be as close as possible to $g$; if both $x$ and $y$ shift are zero (the home node), the returning ant stops;

(R2) if the cell suggested by the DRS (minimization of $x$ and/or $y$ shift, while maintaining $g$) cannot be reached because of a communication failure (an obstacle), the ant selects an obstacle-avoiding move according to the transitional probabilities (see Equation 6); upon this selection the ant keeps to the chosen path until the obstacle is avoided, as recognized by comparison of current $y/x$ ratio with $g$;

(R3) each cycle, a returning ant deposits pheromone in a quantity inversely proportional to the traversed return distance $q$ ($q$ is incremented by 1 each cycle); the deposited pheromone is limited by a predefined maximum $\varphi_{\max}$.

The pheromone is deposited on the cells themselves rather than communication links—we deal with pheromone trail intensities $\varphi_j$ at the cell $j$. While calculating the ant-routing tables, we set $\alpha = 1$ and $\beta = 0$, resulting in

$$a_{i,j} = \frac{\varphi_j}{\sum_{l \in N_i} \varphi_l} \quad \forall j \in N_i.$$

The visibility factor is discarded—to some degree its role is taken by the DRS. In order to update its ant-routing table the cell $i$ needs to receive individual pheromone updates from its neighbors, and this data may be included in the cyclical "OK" messages. In order not to retrace its steps, an ant $k$ located at cell $i$ remembers (for one cycle only) the node $m^k$ where it was at the previous cycle, and excludes it from the feasible neighborhood $N_i^k = N_i \backslash \{m^k\}$. The transition probabilities used on the "follow-trail" step (E3) are calculated according to Equation 5. The returning ants, following the DRS, use transition probabilities only when they face a new obstacle, at the step (R2):

$$p_{i,j}^k = \begin{cases} \frac{a_{i,j}}{\sum_{l \in N_i^k} a_{i,l}} & \text{if} \quad \exists l \in N_i^k, \quad \varphi_l > \overline{\varphi}, \\ 1/N & \text{if} \quad \forall l \in N_i^k, \quad \varphi_l \leq \overline{\varphi}, \\ 1 & \text{if} \quad N_i^k = \emptyset, \quad j = m^k, \end{cases} \tag{6}$$

where $N$ is the number of nodes in the feasible neighborhood $N_i^k$, and $\overline{\varphi}$ is a simple obstacle threshold encouraging more exploration in probabilistic decisions earlier in the process.

The algorithm performs well in dealing with two well-known problems: the blocking problem and the shortcut problem [21]. *Blocking* occurs when a trail that was found by the ants is no longer available due to obstacle(s) and an alternative trail is needed. This is solved by the ants that go along a shorter return path depositing more pheromone than the ants that selected a longer path around an obstacle—simply because the deposited quantity is inversely proportional to the traversed distance. A larger quantity of pheromone attracts more ants. Eventually, a trail is established between a pair of impact nodes. It is important to realize that in any such pair, each impact node generates exploring ants that potentially find the other impact node and return. So the shortest trail is reinforced by returning ants going in opposite directions. The *shortcut* corresponds to a new shorter trail becoming available due to repaired cells—this problem is alleviated by the DRS. Finally, the algorithm produces minimum spanning trees, resulting in reconfigurable impact networks (Figure 7). In summary, this modified ACO algorithm involves

- one type of nondissipative evaporating pheromone, limited by $\varphi_{\max}$;
- simple ant-routing tables containing normalized pheromone values only for immediate neighbors;
- one type of ants with private memory limited to four variables (the relative $x$ and $y$ shifts from the home node, the relative index $m^k$ of the cell where the ant was at the previous cycle, and the ratio $g = y/x$ of the found node's coordinates relative to the home node);
- the dead-reckoning scheme;
- the transition-probability model with the obstacle threshold $\overline{\varphi}$.

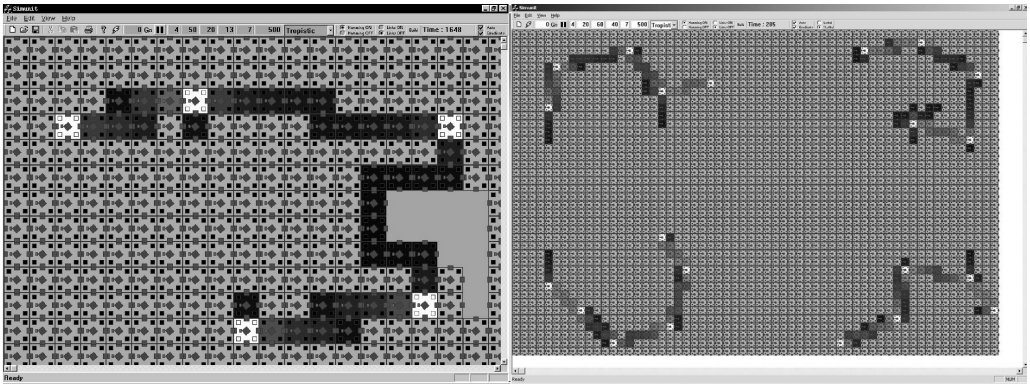

Figure 7. White cells detected non-critical impacts. Left: An example of a reconfigurable impact network. Right: Reconfigurable impact networks without emergent hubs.

The analysis of spatial stability of impact networks is analogous to that conducted for impact boundaries, and is based on the concept of a connected trail fragment (CTF). A CTF is a set $F$ of cells with $\varphi \geq \vartheta$ (where $\vartheta$ is a given threshold), such that every cell in $F$ is connected with at least one other cell in $F$, and there exists no cell outside $F$ that is connected to at least one cell in $F$. The study of the impact network's stability focused on two parameters: the ants' generation interval $T$ and retention rate $\varrho$. As expected, both parameters have critical values corresponding to phase transitions in the connectivity of impact networks. In particular, short intervals ($T \leq 20$) with a high frequency of the ant generation process result in robust, ordered trails; medium-length intervals ($20 < T < 100$) result in highly unstable trails, with the most complex dynamics (trails frequently form and break); and long intervals ($T \geq 100$) generate too few ants to support a trail for any significant time—the ensuing dynamics can be characterized as chaotic. Similar dynamics is observed on varying the retention rates, producing robust trails for the higher retention rates $\varrho \geq 0.97$, unstable trails for the medium rates around $\varrho = 0.93$, and chaotic trails for the lower rates $\varrho < 0.9$, where pheromones evaporate too quickly.

Robust impact networks employ second-level dynamic trails between impact nodes, and exemplify a third-level structure emerging as a result of multi-agent interactions. They have a unique third-order property: a spanning tree topology.

## 7 Conclusions and Future Work

A scale-free network is a web without a spider.—Albert-Laszlo Barabasi [32]

In this article we considered emergence of sensor-communication hierarchies on multiple levels. A first-order unit was chosen to be a cell, followed by self-organization of second-order impact boundaries and third-order impact-surrounding regions, as well as third-order impact networks. Each level of the hierarchy was shown to have distinct higher-order emergent properties, desirable in self-monitoring and self-repairing AAVs. Another important question addressed in this work is how to retain desirable emergent behavior while avoiding potentially adverse patterns of agents' interaction within and between various hierarchical levels. In particular, both cells and communication messages have been shown to need memory (hysteresis): in particular, the communication history length $\mu$, the recovery-delay period $\pi$, the pheromone retention rate $\varrho$ for cells, and time-to-live counters for communication packets. The cell's shutdown ability is related to memory as well, and would help even in the face of error-prone processing in the cell's logic unit. In general, the hysteresis helps to stabilize emergent hierarchies, and a number of quantitative

techniques for measuring spatiotemporal robustness of self-organizing boundaries and networks have been presented. In particular, we focused on identifying the edge of chaos, leading to discovery of evident phase transitions in sensor-communication networks. The presented metrics can be used as fitness functions in evolving the desired self-organizing patterns, and work in this direction is in progress [26]. Ideally, the evolved system should not spend a significant time in either the chaotic or the complex regime; therefore, the phases' identification would better describe the selection pressures on critical hysteretic parameters.

The final system is likely to have bidirectional interactions between all levels, that is, the system will provide actuation as well as sensing functions. While that is not implemented at this stage, the design allows for it. Some plausible functions are excitation of sensors for active measurements, generation of sensor calibration signals, transmission of wireless signals, secondary inspections with roaming virtual or physical agents and/or swarms, and even initiation of some types of repair.

One challenging direction planned for the future work is development of hierarchies supporting coordinated actions over spatially separated areas. In particular, it may be possible to connect together multiple emergent impact spanning trees, creating a Barabasi-Albert *scale-free* network [3]. The majority of nodes in a scale-free network will have one, two, or three links, being part of a specific impact spanning tree, but a few nodes (hubs) will have a large number of links connecting to other impact spanning trees—guaranteeing that the system becomes fully connected. In other words, these hubs may form another level (the fourth) of the main hierarchy, producing fourth-order emergent properties. Figure 7 (right) shows a number of third-order impact spanning trees: there are no emergent hubs, symptomatic of a scale-free network, precisely because our ant algorithm favors formation of shortest paths, necessarily creating disjoint robust clusters. In other words, another algorithm is needed to create a fourth-order scale-free network.

A few well-known emergent properties of scale-free networks will be highly desirable: full connectivity, searchability, and having indegrees and outdegrees of the nodes following power laws. The nodes with the largest numbers of links in a scale-free network play a more important role (e.g., they may coordinate a secondary inspection and/or a repair), while clusters may store some important local information (e.g., impact spread). In general, a reconfigurable scale-free network, emerging as a new hierarchical level, will guide secondary inspections with active sensors as well as mobile virtual agents (task-oriented communication packets, in reality) and even mobile physical agents roaming on the AAV skin. By "guiding" we mean here not the use of a remote control or navigation system, but rather identification of points and regions of interest, and types and schedules of inspections.

Mobile agents may form, we believe, yet another (fifth) level of the main hierarchy—adaptive swarms of inspectors that self-organize *in response* to behavior of the emergent scale-free network. This topic deserves a separate study, of course. We would like to point out, however, a few promising results reported in the recent literature. A quite sophisticated pattern formation model for a class of self-assembling modular robotic units is described by Trianni et al. [24]. Their work lies at the intersection of collective robotics and metamorphic robotics, and also draws from biology: the emergent chaining swarm behavior is analogous to the process of polymerization, while the emergent clustering behavior is similar to the autocatalytic process used by pheromone-depositing bark beetle larvae.

Ideally, there will be a number of parallel feedback loops from global inspection results into the hierarchy itself. This will lead to reconfiguration of the fully connected scale-free sensor-communication-inspection-repair network in such a way that the structure of subsequent actions (topology, schedule, hysteresis, etc.) will change. The overall objective of this feedback is to balance the risk associated with local inspections and responses and the cost of global inspections and repairs.

## Acknowledgments

A part of this work was carried out under a CSIRO contract to NASA. It is a pleasure to record our appreciation to Dr. Ed Generazio (NASA Langley Research Center) for his encouragement and support. The authors are grateful to other members of the AAV project, GREM$^{LAB}$, and the

discussion group on "Entropy and self-organization in multi-agent systems," especially Vadim Gerasimov, Mark Hedley, Nigel Hoschke, Geoff James, Mark Johnson, Chris Lewis, Geoff Poulton, and Andrew Scott, as well as Mark Dras. Mark Foreman's 2002–2003 summer vacation scholarship was provided by CSIRO Telecommunication and Industrial Physics. We would also like to thank the five anonymous referees for their careful reading and several insightful comments that led to a substantial improvement of this article.

## References

1. Abbott, D., Doyle, B., Dunlop, J., Farmer, A., Hedley, M., Herrmann, J., James, G., Johnson, M., Joshi, B., Poulton, G., Price, D., Prokopenko, M., Reda, T., Rees, D., Scott, A., Valencia, P., Ward, D., & Winter, J. (2002). *Development and evaluation of sensor concepts for ageless aerospace vehicles. Development of concepts for an intelligent sensing system* (Technical report NASA/CR-2002-211773). Hampton, VA: NASA Langley Research Center.

2. Abbott, D., Batten, A., Carpenter, D. C., Dunlop, J. B., Edwards, G. C., Farmer, A. J., Gaffney, B., Hedley, M., Hoschke, N., Isaacs, P., Johnson, M. E., Lewis, C. J., Murdoch, A., Poulton, G. T., Price, D. C., Prokopenko, M., Rees, D. E., Scott, D. A., Seneviratne, S., Valencia, P., Wang, P., & Whitnall, D. F. (2003). *Development and evaluation of sensor concepts for ageless aerospace vehicles. Report 4: Phase 1 — implementation of the concept demonstrator* (Confidential report TIPP 1898). Lindfield, Australia: CSIRO Telecommunications and Industrial Physics.

3. Barabasi, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*, 509–512.

4. Colorni, A., Dorigo, M., & Maniezzo, V. (1992). Distributed optimization by ant colonies. In F. J. Varela & P. Bourgine (Eds.), *Toward a practice of autonomous systems: Proceedings of the 1st European Conference on Artificial Life, Paris, December 1991* (pp. 134–142). Cambridge, MA: MIT Press.

5. Dorigo, M., Maniezzo, V., & Colorni, A. (1991). *The Ant System: An autocatalytic optimizing process* (Technical report 91-016 revised). Milan, Italy: Politecnico di Milano.

6. Dorigo, M., Maniezzo, V., & Colorni, A. (1996). The Ant System: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics, B*, *26*(1), 1–13.

7. Dorigo, M., & Di Caro, G. (1999). Ant colony optimization: A new meta-heuristic. In *Proceedings of 1999 Congress on Evolutionary Computation, Washington, DC* (pp. 1470–1477). Piscataway, NJ: IEEE Press.

8. Durbeck, L. J. K., & Macias, N. J. (2002). Defect-tolerant, fine-grained parallel testing of a cell matrix. In J. Schewel, P. B. James-Roxby, H. H. Schmit, & J. T. McHenry (Eds.), *Proceedings of SPIE ITCom 2002 Series 4867* (pp. 71–85). Boston, MA: SPIE.

9. Erdös, P., & Renyi, A. (1961). On the strength of connectedness of random graphs. *Acta Mathematica Scientia Hungarica*, *12*, 261–267.

10. Estrin, D., Govindan, R., Heidemann, J., & Kumar, S. (1999). Next century challenges: Scalable coordination in sensor networks. In H. Kodesh, V. Bahl, T. Imielinski, & M. Steenstrup (Eds.), *Proceedings of the 5th Annual ACM/IEEE International Conference on Mobile Computing and Networks, Seattle* (pp. 263–270). New York: ACM Press.

11. Foreman, M., Prokopenko, M., & Wang, P. (2003). Phase transitions in self-organising sensor networks. In W. Banzhaf, T. Christaller, P. Dittrich, J. T. Kim, & J. Ziegler (Eds.), *Advances in artificial life—Proceedings of the 7th European Conference on Artificial Life (ECAL), Dortmund* (pp. 781–791). Heidelberg, Germany: Springer-Verlag.

12. Green, D. (1993). Emergent behavior in biological systems. In D. G. Green & T. J. Bossomaier (Eds.), *Complex systems—from biology to computation* (pp. 24–35). Amsterdam: IOS Press.

13. Langton, C. (1991). Computation at the edge of chaos: Phase transitions and emergent computation In S. Forest (Ed.), *Emergent computation* (pp. 12–37). Cambridge, MA: MIT Press.

14. Lovatt, H., Poulton, G., Price, D., Prokopenko, M., Valencia, P., & Wang, P. (2003). Self-organising impact boundaries in ageless aerospace vehicles. In J. S. Rosenschein, T. Sandholm, M. Wooldridge, & M. Yokoo (Eds.), *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multi-agent Systems, Melbourne* (pp. 249–256). New York: ACM Press.

15. Macias, N. J., & Durbeck, L. J. K. (2002). Self-assembling circuits with autonomous fault handling. In A. Stoica, J. Lohn, R. Katz, D. Keymeulen, & R. S. Zebulum (Eds.), *Proceedings of the 4th NASA/DoD Conference on Evolvable Hardware, Alexandria* (pp. 46–55). Los Alamitos, CA: IEEE Computer Society Press.

16. Mjolsness, E., & Tavormina, A. (2000). The synergy of biology, intelligent systems, and space exploration. *IEEE Intelligent Systems—AI in Space*, *3/4*, 20–25.

17. Nagpal, R. (2002). Programmable self-assembly using biologically-inspired multiagent control. In C. Castelfranchi & W. L. Johnson (Eds.), *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multi-agent Systems, Bologna* (pp. 418–425). New York: ACM Press.

18. Polani, D. (2003). Measuring self-organization via observers. In W. Banzhaf, T. Christaller, P. Dittrich, J. T. Kim, & J. Ziegler (Eds.), *Advances in artificial life—Proceedings of the 7th European Conference on Artificial Life (ECAL), Dortmund* (pp. 667–675). Heidelberg: Springer-Verlag.

19. Price, D., Scott, A., Edwards, G., Batten, A., Farmer, A., Hedley, M., Johnson, M., Lewis, C., Poulton, G., Prokopenko, M., Valencia, P., & Wang, P. (2003). An integrated health monitoring system for an ageless aerospace vehicle. In F.-K. Chang (Ed.), *Proceedings of the 4th International Workshop on Structural Health Monitoring, Stanford, CA* (pp. 310–318). Lancaster, PA: DEStech Publications.

20. Rasmussen S., Baas N. A., Mayer, B., Nilsson, M., & Olesen M. W. (2001). *Ansatz* for dynamical hierarchies. *Artificial Life*, *7*(4), 329–353.

21. Schoonderwoerd, R., Holland, O. E., Brutten, J. L., & Rothkrantz, L. J. M. (1997). Ant-based load balancing in telecommunication networks. *Adaptive Behaviour*, *5*(2), 169–207.

22. Schweitzer, F., & Tilch, B. (2002). Self-assembling of networks in an agent-based model. *Physical Review E, 66*, 026113, 1–9.

23. Solé, R. V., & Valverde, S. (2004). Information theory of complex networks: On evolution and architectural constraints. In E. Ben-Naim, H. Frauenfelder, & Z. Toroczkai (Eds.), *Complex networks, Lecture Notes in Physics, Vol. 650* (pp. 189–210). Berlin: Springer-Verlag.

24. Trianni, V., Labella, T. H., Gross, R., *S*ahin, E., Dorigo, M., & Deneubourg, J.-L. (2002). *Modeling pattern formation in a swarm of self-assembling robots* (Technical report TR/IRIDIA/2002-12). Brussels, Belgium: IRIDIA.

25. Wang, P., Valencia, P., Prokopenko, M., Price, D., & Poulton, G. (2003). Self-reconfigurable sensor networks in ageless aerospace vehicles. In U. Nunes, A. T. de Almeida, A. K. Bejczy, K. Kosuge, & J. A. T. Machado (Eds.), *IEEE proceedings of the 11th International Conference on Advanced Robotics (ICAR), Coimbra* (pp. 1098–1103). Coimbra, Portugal: Universidade de Coimbra.

26. Wang, P., & Prokopenko, M. (2004). Evolvable recovery membranes in self-monitoring aerospace vehicles. In S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, & J.-A. Meyer (Eds.), *From animals to animats VIII—proceedings of the 8th International Conference on the Simulation of Adaptive Behaviour, Los Angeles* (pp. 509–518). Cambridge, MA: MIT Press.

27. Wehner, R., & Srinivasan, M. V. (1981). Searching behaviour of desert ants, genus Cataglyphis (Formicidae Hymenoptera). *Journal of Computational Physiology*, *142*, 315–338.

28. Wuensche, A. (1999). Classifying cellular automata automatically: Finding gliders, filtering, and relating space-time patterns, attractor basins, and the *Z* parameter. *Complexity*, *4*(3), 47–66.

29. Dillard, A. (1974). *Pilgrim at Tinker Creek* (Chapter 7). New York: Harper Row.

30. Mackay, C. (1932). *Extraordinary popular delusions and the madness of crowds, preface to edition of 1852: Memoirs of extraordinary popular delusions and the madness of crowds* (p. xx). New York: Farrar, Strauss and Giroux.

31. Burckhardt, J. (1979). *Reflections on history* (Chapter 4). Indianapolis: Liberty Classics. (Translation of *Weltgeschichtliche Betrachtungen*, 1905.)

32. Barabasi, A.-L. (2002). *Linked: The new science of networks* (p. 221). Cambridge, MA: Perseus.