

Phase Transitions in Self-organising Sensor Networks

Mark Foreman¹, Mikhail Prokopenko², Peter Wang²

Ageless Aerospace Vehicles and Smart Spaces projects
Intelligent Interactive Technology

¹ CSIRO Telecommunications and Industrial Physics

² CSIRO Mathematical and Information Sciences

Locked Bag 17, North Ryde, NSW 1670, Australia

{mark.foreman, mikhail.prokopenko, peter.wang}@csiro.au

Abstract. In this paper we consider a multi-cellular sensing and communication network, embedded in an *ageless aerospace vehicle*, that is expected to detect and react to impact location, intensity and damage over a wide range of impact energies. In particular, we investigate self-organisation of *impact boundaries* enclosing critically damaged areas, and measure their spatiotemporal robustness. The presented quantitative information-theoretic techniques clearly identify phase transitions, separating chaotic dynamics from ordered and robust patterns.

1 Introduction

The research results presented in this paper were obtained as part of the joint CSIRO-NASA Ageless Aerospace Vehicle (AAV) project. The aim of the AAV project is to develop and critically examine concepts for integrated smart sensing and communication networks, with the ultimate goal of developing a self-monitoring, self-repairing aerospace vehicle [1, 6, 11].

A modular (multi-cellular) 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 micro-particles to meteoroids. The network is expected to self-organise in the face of damage to its parts so that robust monitoring and reporting continues as long as possible. To achieve this, we modelled and simulated a multi-agent system made of “cells”, that will not only form a physical shell for an aerospace vehicle, but will also have sensors, logic, and communications. The primary principle that is followed in our work is the *emergence of global response* as a result of interactions involving transfer of information embedded locally — in order to completely avoid or reduce the number of single points-of-failure. In other words, without centralised controllers, agents (cells) are expected to self-organise and survive on the basis of local, rather than global, information (no single agent has access to what everyone else is doing). Single cells may need to make fast and automatic responses to sudden damage, while collections of cells may solve more complex tasks, for example, produce an impact boundary with desired characteristics [6] or form a spanning tree connecting cells that detected non-critical impacts [11]. Importantly, the behaviour 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/replace it.

Recent advances in sensor networks and micro-electro-mechanical devices led to the idea of localised algorithms, in which simple local node behaviours achieve a desired global objective [3, 7], while communicating only with nodes within some neighbourhood. However, despite some progress, there is a lack of a unifying methodology underlying design of localised algorithms. The main question is how to produce and retain desirable emergent behaviour while avoiding potentially adverse patterns of agents' interaction. Some promising results have been reported by Nagpal [8], in context of amorphous computing, where programmable self-assembly was demonstrated using biologically-inspired multi-agent control. Nagpal defined a small class of primitives (eg., gradients, neighbourhood query, polarity inversion, cell-to-cell contact, etc.), a set of global operations, and a translation implementing the global operations as localised algorithms using the set of primitives. However, the question of how to obtain and inter-connect global operations themselves is left unanswered. Moreover, the task of discovering the correct primitives and translating them into localised agent programs would have to be repeated for each domain.

The problem of global response engineering in multi-agent networks motivates our search for patterns and invariants in self-organisation of multi-agent systems monitoring and repairing impacts on the AAV. Our proposed methodology is based on an iterative process including the following steps: a) forward simulation leading to emergent behaviour (for a class of localised algorithms dealing with impacts of various strengths); b) quantitative measurement of bounded emergent behaviour (using information-theoretic metrics for phase transitions); c) evolutionary modelling of the desired global emergent behaviour, where the fitness functions correspond to the metrics obtained at step b). The work on impact boundaries [6] and impact networks [11] modelled the step a). This paper studies the step b).

2 Impact Boundaries

Ideally, a modular multi-cellular AAV skin should trace impact boundaries and spread despite connectivity disruptions and cell failures — analogous to the clotting of a wound on a mammal and the regeneration of neurons by re-growing severed axons within a myelin sheath. One of the immediate tasks is the formation of impact boundaries enclosing critically damaged areas. It is highly desirable that such boundaries form continuously connected closed circuits, and are robust to fluctuations caused by proximity to the impact. In short, the aim is to achieve spatiotemporal stability in impact boundaries.

Let us briefly sketch here the AAV simulator and the algorithms producing self-organising continuous and closed impact boundaries that are presented in more detail in [6]. In the AAV simulator events may occur in parallel, so the simulation is, essentially, a state machine that sweeps through the cells, updating their current state on a regular basis. Cells are represented as objects (polygons) on a two-dimensional plane, where they interact *only* with their immediate neighbours in von Neumann neighbourhood, through connected (geometrically overlapping) communication ports. Furthermore, the AAV simulator has the ability to simulate simple environmental effects, such as the incidence of small impacts.

An important characteristic of a critical impact is that, typically, energy of the impact dissipates throughout some neighbourhood, destroying cells close to the point of

the impact (epicenter) and damaging the communication links between the survived cells. The communication damage is simulated by assigning a probability of a bit error dependent on proximity of the affected communication port to the epicenter. The proximity dependency underlying the probability distribution can be modelled in many ways, and we investigated a range of functions — from linear to exponential decreases.

Our first target was formation of a boundary that is not necessarily continuously connected as a closed circuit, but is well placed in separating the cells that suffered unrecoverable communication damage (including those that were completely destroyed) from the cells that are able to communicate to their normal functional capacity. The initial algorithm enables self-organisation of both an internal “scaffolding” and the “frame” of the desired closed boundary, regardless of cells shape (triangular or square), and includes the following communication policies and behaviours:

- (i) every cycle each cell sends an “OK” message to all its neighbours;
- (ii) upon receiving the “OK” message each cell replies with an “Acknowledgement” message;
- (iii) if all neighbour cells failed to communicate the “OK” message, switch to the scaffolding state S_s ;
- (iv) if at least one neighbour cell failed to communicate the “OK” message and at least one neighbour cell did communicate the “OK” message, switch to the frame boundary state S_f .
- (v) if all neighbour cells failed to communicate the “OK” message or there is no “Acknowledgement” message from any neighbour cell, stop sending messages.

In order to produce *continuously* linked closed boundary, we need the following:

- (vi) if the cell state is S_f , and there are at least two communicating neighbours β_1 and β_2 , switch to the closed state S_c .
- (vii) if the cell state is S_c , 1) determine a (the) cell α that failed to communicate; 2) determine two communicating neighbour cells β_1 and β_2 nearest to the cell α on opposite sides (clockwise and counter-clockwise relative to α); 3) establish a boundary link between ports to β_1 and β_2 ; 4) map the directions to α , β_1 and β_2 to a desired direction γ ; and 5) send a “Connect(τ, γ)” message to these neighbours with a “time to live” parameter τ ;
- (viii) upon receiving “Connect(τ, γ)” message from the cell α , if the cell state is not S_c , 1) switch to the state S_c ; 2) determine the communicating neighbour cell β nearest to the cell α in the direction γ ; 3) establish a boundary link between ports to α and β ; 4) if $\tau > 0$, map the directions to α and β to a desired direction γ' and 5) send a “Connect($\tau - 1, \gamma'$)” message to the neighbour β .

The “time to live” parameter τ 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. In general, the described policy achieves the desired robustness and continuity of self-organising impact boundaries for a variety of cell shapes, impact energy dissipation profiles and communication damage probability models (Figure 1). However, on some occasions the boundary is erratic and unstable — some cells frequently change their states from “frame” to “scaffolding” to normal, etc. The reason for this spatial and temporal instability is that, a simple detection of

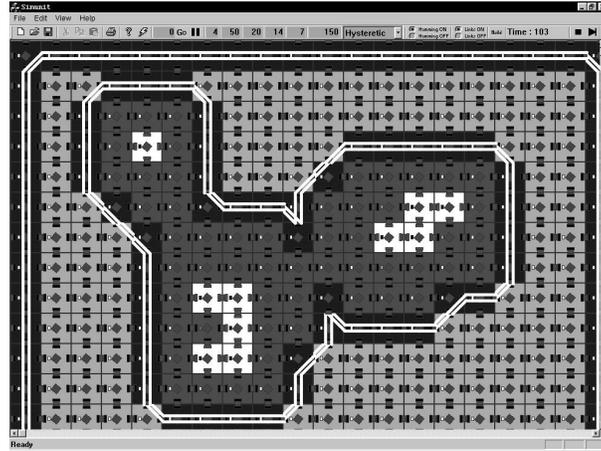


Fig. 1. A stable impact boundary: square cells ($\rho = 10$). White cells are destroyed, dark-grey cells form “scaffolding”, black cells form “frame”. Boundary links are shown as white double-lines.

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. The length ρ of such a filter (i.e., the cell-to-cell communication history, allowed to be kept by each cell), is an important parameter, directly affecting spatial and temporal stability of self-organising impact boundaries. Our experiments verified that larger values of ρ lead to stable boundaries, while $\rho = 0$ (no hysteresis at all) could not produce any boundaries (Figure 2).

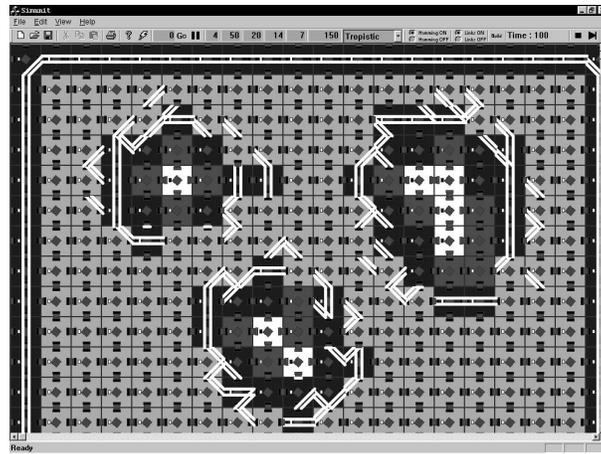


Fig. 2. An unstable impact boundary: square cells with $\rho = 0$.

3 Spatiotemporal Stability and Phase Transitions

The analysis of spatiotemporal stability of self-organising impact boundaries with respect to the parameter ρ is the primary subject of our investigation. In particular, we characterise dynamics of multi-cellular impact boundaries in terms of generic information-theoretic properties, such as the Shannon entropy, and pinpoint certain phase transitions.

Information-theoretic methods are applied in many areas exhibiting multi-agent interactions. For instance, Cellular Automata (CA) are a well-studied class of discrete dynamical systems with emergent behaviour resulting from local and short range interactions, and where information-theoretic measures of complexity (such as Shannon entropy of certain frequency distributions) were effectively used to categorise and classify distinct emergent configurations and phase transitions between them [13, 5]. Langton has shown in his seminal work [5] that an increase in the mutual information (defined as a function of individual cell entropies for a particular value of the λ parameter) is an indication of a phase transition from “order” to “chaos”. Wuensche [13] has used a similar quantitative metric — variance of input-entropy over time — in classifying rule-space of 1-dimensional CAs into ordered, complex and chaotic cases, related to Wolframs’s qualitative classes of CA behaviour [12]. It has also been pointed out by Suzudo [10] that the entropy trajectory (plotting spatial entropy against temporal entropy) is a useful descriptor for a variety of self-organised patterns. For example, certain (non-complex) types of CA have simple entropy trajectory and do not renew their self-organised spatiotemporal pattern. Complex CA have an irregular entropy trajectory in the temporal-spatial entropy plane (similar to Figure 7), and renew the pattern.

The information-theoretic analysis of phase transitions is typically based on the notion of *mutual information*. For instance, Langton [5] investigated the mutual information of CA, defined as a simple function of the individual cell entropies, $H(A)$ and $H(B)$, and the entropy of the two cells considered as a joint process, $H(A, B)$:

$$I(A; B) = H(A) + H(B) - H(A, B)$$

and related it to phase transitions. In particular, trajectories of entropy $H(A)$ and mutual information $I(A; B)$ between a cell and itself at the next time-step were obtained while varying the parameter λ — the ratio of cells with a given property (“live” cells, for example). Interestingly, the individual cell entropy $H(A)$ was increasing with λ , and showed a discrete jump between low and high entropy values. This evidence pointed to a first-order phase transition, similar to that observed between the solid and fluid phases of matter, however the fact that the gap is not completely empty suggested the possibility of second-order transition [5]. Another intriguing feature was that the average mutual information has a distinct peak at the transition point:

the average mutual information is essentially zero before the transition point, it jumps to a moderate value at the transition, and then decays slowly with increasing λ . The jump in the mutual information clearly indicates the onset of the chaotic regime, and the decaying tail indicates the approach to effectively random dynamics.

It is important to realise that there are two opposing forces shaping the trajectory of mutual information — individual diversity and interdependence. In the case studied by Langton, increasing λ led to an increase in diversity and growth of individual entropy. At the same time, correlation and interdependence (reflected in the entropy of the joint process) were decreasing: at the start the cells were overly dependent, and at the end they were practically independent of each other. The “battle” between these two forces is the most intense at the transition point, or in other words, at the *edge of chaos*, where the system dynamics exhibits the most complexity.

A metric based on mutual information is suitable precisely because it incorporates these two tendencies into a single expression. Another way to achieve this combination is to use *a variance of input-entropy*. Wuensche [13] characterised rule-spaces of 1-dimensional cellular automata with the Shannon entropy of rules’ frequency distribution. More precisely, given a rule-table (the rules that define a CA), the input-entropy at time step t is defined as

$$S^t = - \sum_{i=1}^m \frac{Q_i^t}{n} \log \frac{Q_i^t}{n},$$

where m is the number of rules, n is the number of cells (system size), and Q_i^t is the look-up frequency of rule i at time t — the number of times this rule was used at t across the CA. 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 follows an 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. This variance, in a sense, captures mutual information contained in CA dynamics, combining both individual rules’ diversity and their interdependence (correlation). 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 input-entropy variance points to a phase transition again.

We believe that a metric modelled after the input-entropy S^t may capture only the temporal stability of impact boundaries, but not the spatial stability of continuous and closed circuits. Put simply, the input-entropy S^t may be very low and may vary only slightly for temporally stable patterns that keep their disconnected shape. So it would have to be complemented with another metric designed for spatially connected patterns. The random graphs theory that dates back to seminal works of Erdos and Renyi [2] suggests to use the size of largest connected sub-graph (LCS) and its variance or standard deviation obtained across an ensemble of graphs, as an indicator of phase transitions in the graph connectivity. It is well-known that critical changes occur in connectivity of a directed graph 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 [4]. A metric based on the variance in the size of the LCS may capture connectivity in impact boundaries. Again, by itself, it would be insufficient: a continuous boundary may change its shape over time, without breaking into fragments, and keeping the size of LCS almost constant. Together, temporal and spatial metrics may capture both important aspects — steadiness and continuity of impact boundaries.

4 Experimental Results

In order to analyse temporal stability, we consider state changes in each cell at every time step. There are 6 symmetric boundary links possible in each square cell, connecting ports “left-right”, “top-bottom”, “left-top”, etc. (Figure 1). Thus, there are 2^6 possible boundary states (including “no-boundary”), and $m = 2^{12}$ transitions. These transitions describe space-time dynamics inside an 1-cell neighbourhood (in CA terms), and exemplify a more general case of p-cell neighbourhoods (von Neumann neighbourhood

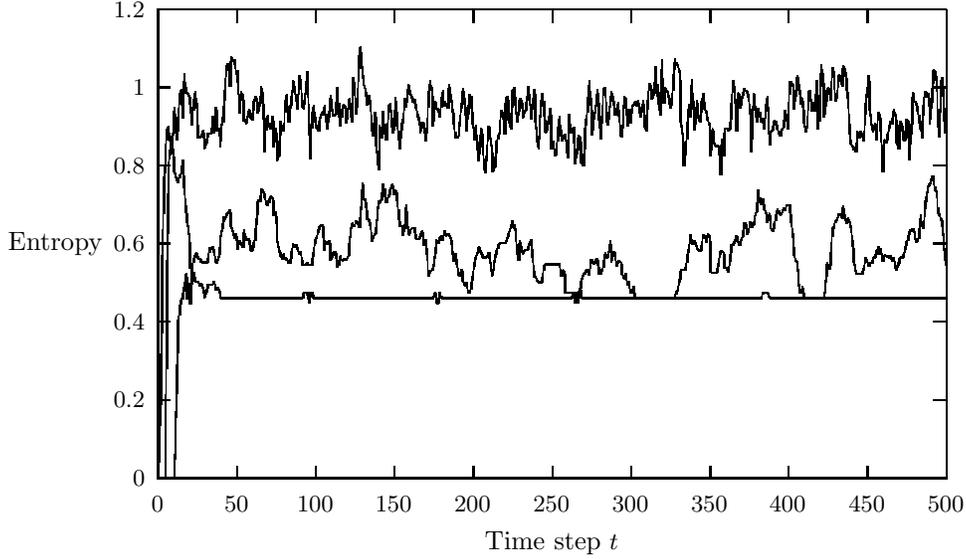


Fig. 3. Temporal entropy $H_{temp}(\rho)$ for $\rho = 0$ (top), $\rho = 4$ (middle), and $\rho = 10$ (bottom).

with $p = 5$; Moore neighbourhood with $p = 9$, etc.). For each value of the parameter ρ (varying from 0 to 20), we calculate the entropy $H_{temp}(\rho)$ of a particular frequency distribution $S_i^t(\rho)$, 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 [13], we define input-entropy as

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

where n is the system size (the total number of cells), and $S_i^t(\rho)$ is the look-up frequency of the transition i at time t . Each experiment simulates an impact at a predefined cell, and lasts 500 cycles; the first 2ρ cycles are excluded from the distribution $S_i^t(\rho)$ in order not to penalise longer history (filter) lengths. Figure 3 plots temporal entropy $H_{temp}(\rho)$ for different history lengths ρ . The $\rho = 10$ plot is almost a straight line — there is practically no variation in the entropy values, as the boundary is stable. We conducted three experiments for every value of ρ , calculating standard deviation $\sigma_{temp}(\rho)$. The plot of $\sigma_{temp}(\rho)$ is shown on Figure 4, clearly identifying a phase transition for some critical value of ρ between $\rho = 4$ and $\rho = 5$.

The analysis of 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 exist no cell outside F which is connected to at least one cell in F (an analogue of a maximally connected subgraph or a graph component). We carried out 63 more experiments (3 for each value of ρ between 0 and 20), calculating the average size $H_{sp}(\rho)$ of CBF's in self-organising impact boundaries at each time-point, and its variance $\sigma_{sp}(\rho)$ over time. Figure 5 plots $H_{sp}(\rho)$ for $\rho = 0$ (bottom), $\rho = 4$ (strongly fluctuating), and $\rho = 10$ (almost straight line). It is clear that the impact boundary in this case has a size 24, and the hysteretic behaviour of length $\rho = 10$ leads to stable self-organisation. On the other hand, the cells without any hysteresis ($\rho = 0$) are not capable of self-

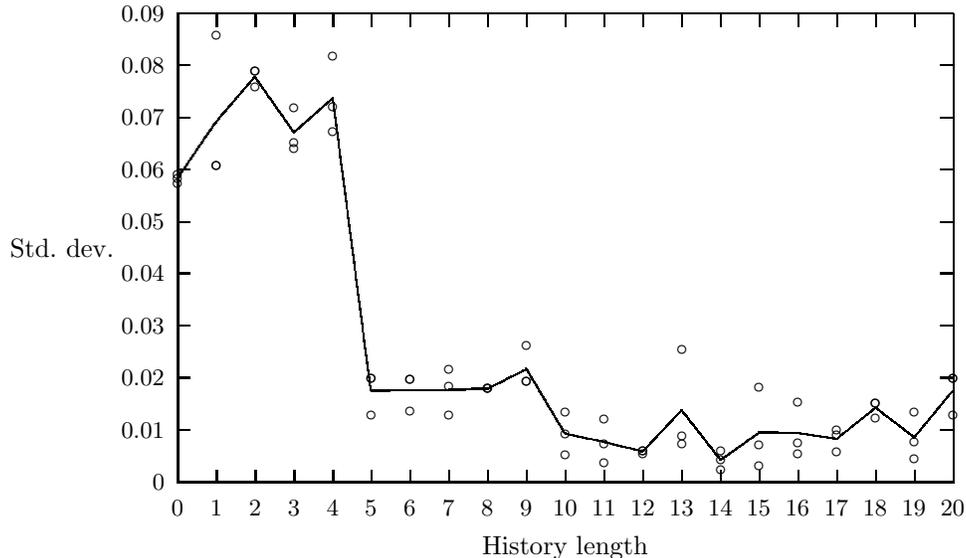


Fig. 4. Standard deviation $\sigma_{temp}(\rho)$ of temporal entropy $H_{temp}(\rho)$.

organisation at all — an average CBF contains about 5 cells. The intermediate case ($\rho = 4$) 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 behaviour leads to “order”, while tropistic behaviour, using a direct mapping from communication inputs to actions ($\rho = 0$), produces “chaotic” dynamics. Somewhere in the middle, we observe “complex” dynamics, and Figure 6 makes the phase transition apparent — again, between $\rho = 4$ and $\rho = 5$. The entropy trajectory in the temporal-vs-spatial plane is shown in Figure 7, plotting $\sigma_{sp}(\rho)$ against $\sigma_{temp}(\rho)$. The bottom-left “order” corner (the points after the phase transition) is clearly contrasted with the top-right “chaos” (the points before the phase transition). The small cluster in the bottom-centre region corresponds to tropistic behaviour ($\rho = 0$). Absence of any points in the middle indicates a first-order phase transition.

5 Conclusions

In this paper, we considered formation of impact boundaries that enclose critically damaged areas on a multi-cellular skin of ageless aerial vehicles. A number of quantitative techniques measuring spatiotemporal robustness of self-organising boundaries was presented. In particular, we focussed on identifying the “edge of chaos”, leading to discovery of evident phase transitions in sensor networks. In future, we intend to use the presented metrics as fitness functions — in *evolving the desired global behaviour*.

Acknowledgements A part of this work was carried out under 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, especially Philip Valencia, Don Price and Geoff Poulton. Mark Foreman’s summer vacation scholarship was provided by CSIRO Telecommunication and Industrial Physics, and we would like to thank Geoff James in particular. The authors also acknowledge the helpful comments of the reviewers.

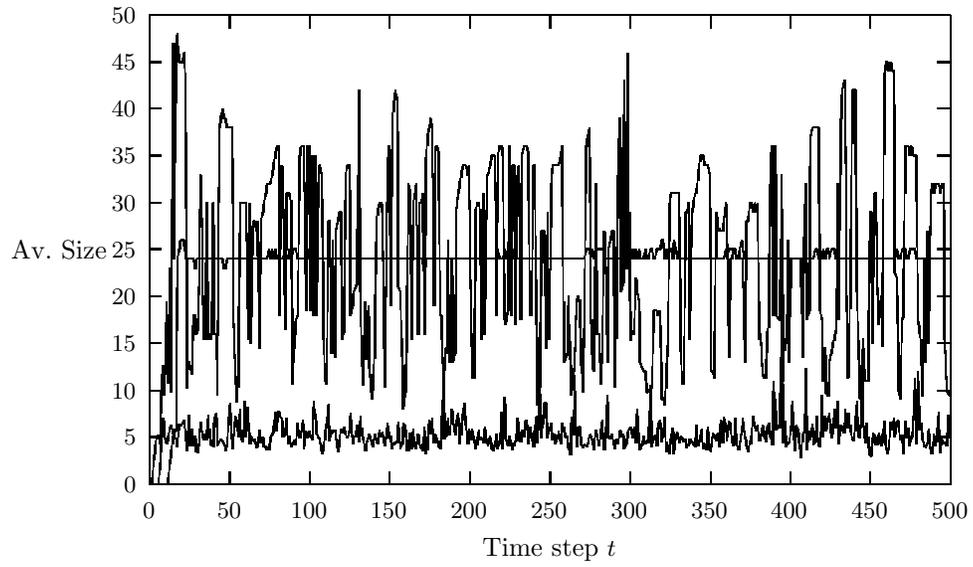


Fig. 5. Average size $H_{sp}(\rho)$ of CBF's, for $\rho = 0$ (bottom), $\rho = 4$ (strongly fluctuating), and $\rho = 10$ (almost straight line).

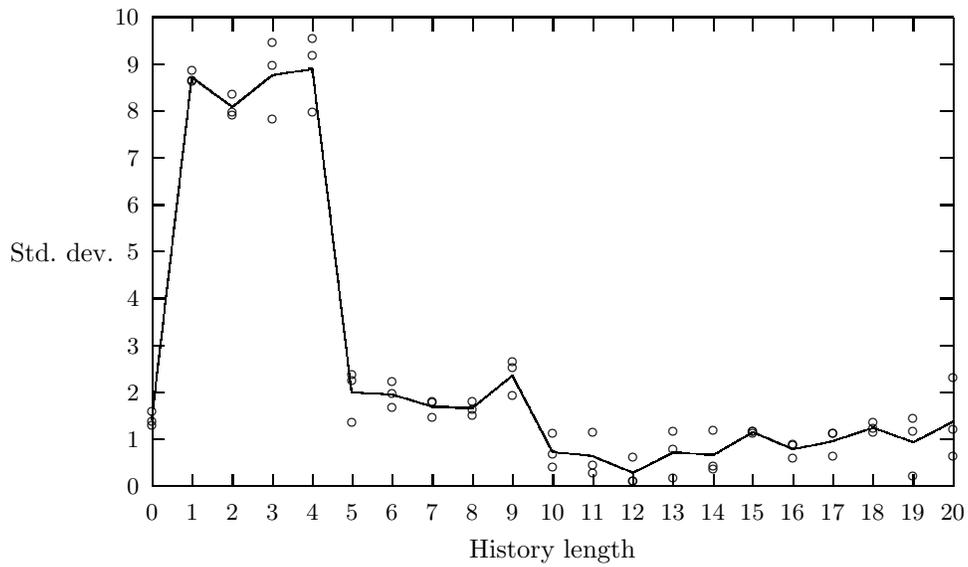


Fig. 6. Standard deviation $\sigma_{sp}(\rho)$ of average size $H_{sp}(\rho)$.

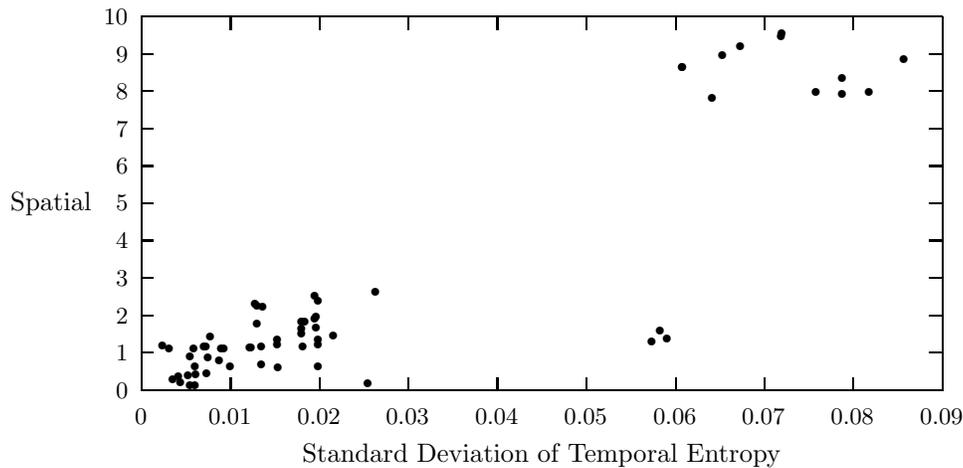


Fig. 7. The entropy trajectory: $\sigma_{sp}(\rho)$ against $\sigma_{temp}(\rho)$.

References

1. David Abbott, Briony Doyle, John Dunlop, Tony Farmer, Mark Hedley, Jan Herrmann, Geoff James, Mark Johnson, Bhautik Joshi, Geoff Poulton, Don Price, Mikhail Prokopenko, Torsten Reda, David Rees, Andrew Scott, Philip Valencia, Damon Ward, and John Winter. Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles. Development of Concepts for an Intelligent Sensing System. NASA technical report NASA/CR-2002-211773, Langley Research Center, Hampton, Virginia.
2. Paul Erdos and Alfred Renyi. On the strength of connectedness of random graphs. *Acta Mathematica Scientia Hungary* 12, 261-267, 1961.
3. Deborah Estrin, Ramesh Govindan, John Heidemann, Satish Kumar. Next Century Challenges: Scalable Coordination in Sensor Networks. In Proceedings of the 5th Annual International Conference on Mobile Computing and Networks (MobiCOM '99), Seattle, 1999.
4. David Green. Emergent Behavior in Biological Systems. *Complexity International*, 1, 1994.
5. Christopher Langton. Computation at the Edge of Chaos: Phase Transitions and Emergent Computation. In S. Forest (ed.), *Emergent Computation*, MIT, 1991.
6. Howard Lovatt, Geoff Poulton, Don Price, Mikhail Prokopenko, Phil Valencia and Peter Wang. Self-organising Impact Boundaries in Ageless Aerospace Vehicles. To appear in Proc. of the 2nd Int'l Joint Conf. on Autonomous Agents and Multi-Agent Systems, 2003.
7. Nicholas J. Macias and Lisa J.K. Durbeck. Self-assembling circuits with autonomous fault handling. In Proceedings of NASA/DoD Conference on Evolvable Hardware, 2002.
8. Radhika Nagpal. Programmable self-assembly using biologically-inspired multiagent control. In the Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'02), Bologna, Italy, July 2002.
9. Claude E. Shannon and Warren Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, 1949.
10. Tomoaki Suzudo. The entropy trajectory: A perspective to classify complex systems. In Proceedings of the Int'l Symposium on Frontiers of Time Series Modeling, Tokyo, 2000.
11. Peter Wang, Philip Valencia, Mikhail Prokopenko, Don Price, and Geoff Poulton. Self-reconfigurable sensor networks in ageless aerospace vehicles. Submitted to the 11th International Conference on Advanced Robotics, ICAR-03, Portugal.
12. Stephen Wolfram. Universality and Complexity in Cellular Automata. *Physica D*, 10, 1984.
13. Andrew Wuensche. Classifying Cellular Automata Automatically: Finding gliders, filtering, and relating space-time patterns, attractor basins, and the Z parameter. *Complexity*, Vol. 4, no. 3, 47-66, 1999.