Evolvable Recovery Membranes in Self-monitoring Aerospace Vehicles

Peter Wang and Mikhail Prokopenko
Centre for Intelligent Systems Design
CSIRO Information and Communication Technologies Centre
Locked bag 17, North Ryde, NSW 1670, Australia
{peter.wang, mikhail.prokopenko}@csiro.au

Abstract

In this paper we proposed and verified a methodology underlying the design of localised algorithms for complex multi-agent systems, exemplified by self-monitoring aerospace vehicles. In particular, we considered the emergence of self-organising impact boundaries and recovery membranes, separating damaged and/or potentially recovering regions from less affected agents. In order to identify phase transitions in system's dynamics, we investigated graph-theoretic and information-theoretic metrics, and incorporated them within fitness functions for a genetic algorithm. The GA involved a generation gap strategy and targeted a response time as well as spatial connectivity, temporal persistence and size of emergent boundaries and membranes. A variety of stable spatiotemporal patterns were produced under selection pressure, highlighting the potential for the *design at the edge of chaos*.

1. Introduction

Robust or "ageless" aerospace vehicles (AAVs) are expected to be capable of structural self-assessment and repair. The research results presented in this paper were obtained as part of the joint CSIRO-NASA AAV project. At this stage, the project is concerned primarily with the self-assessment, or self-monitoring, functions, but it will progress towards the incorporation of damage prognostic models enabling a predictive damage mitigation decision-making process. The initial goal of the present AAV Concept Demonstrator (Abbott et al., 2002, Abbott et al., 2003) is the detection and characterisation of high velocity impacts caused, for instance, by micro-meteoroids in space.

In general, the self-monitoring functions can be divided between those carried out by distributed sensors and decentralised processing and communication on the skin or within the structure, and those that could be more effectively provided by autonomous robotic non-destructive evaluation (NDE) agents deployed to monitor damage or integrity of the vehicle structure. In this paper, we consider the first type of agents embedded in the skin of the AAV. In particular, we shall consider a multi-cellular sensing and communication network. Each cell in this network has the ability to sense via its 4 piezoelectric sensors, process the inputs, and communicate through its 4 communication ports (Abbott et al., 2003), as shown in Figures 1 and 2.

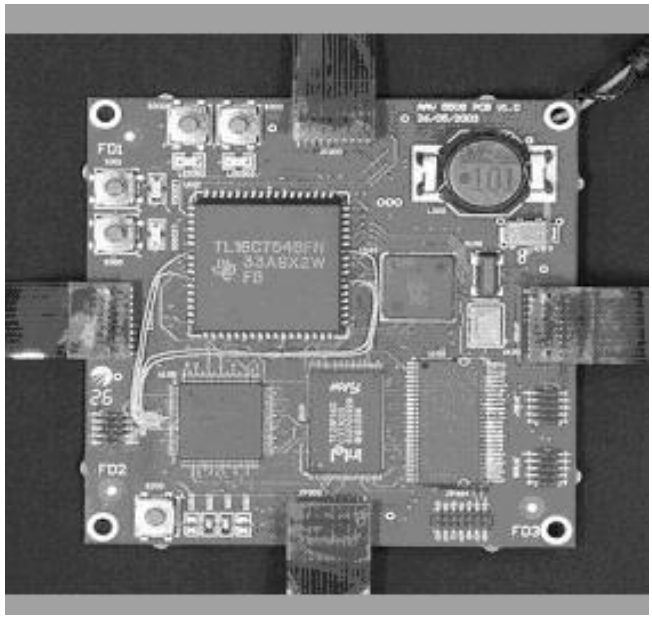

Figure 1: A single square cell with 4 communication ports: Concept Demonstrator view of a printed circuit board.

A two-dimensional array of cells is modelled in a biologically-inspired hybrid Concept Demonstrator, some cells existing in dedicated hardware (a cell per micro-processor) and some residing within interconnected personal computers (a number of cells per PC) (Abbott et al., 2003). We also used a stand-alone Simulator capable of simulating some simple environmental effects such as particle impacts of various energies.

The solution adopted for the problem of handling distributed sensor data and making the self-monitoring

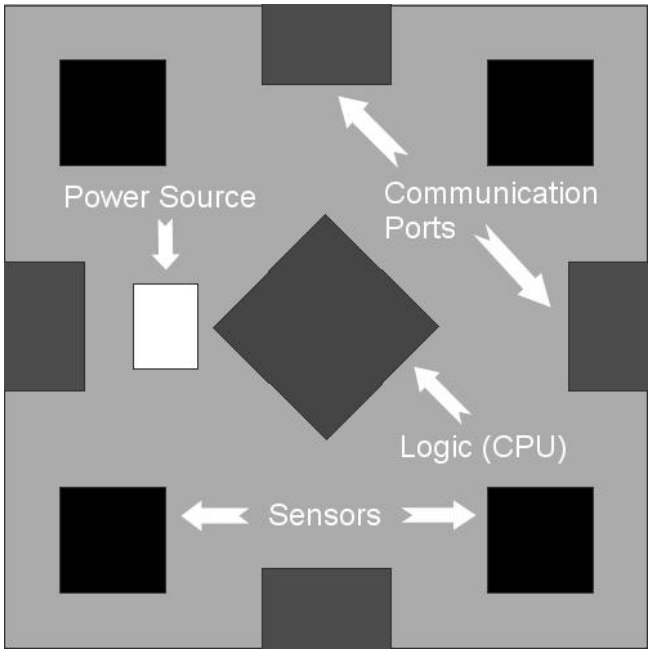

Figure 2: A single square cell with 4 communication ports: AAV Simulator view.

system robust, is to distribute the processing as well, so that most computation takes place near the data source, and there is no single or small number of 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 information about any others but its neighbours). This approach uses the idea of localised algorithms (Estrin et al., 1999, Macias and Durbeck, 2002, Durbeck and Macias, 2002, Prokopenko et al., 2004), in which simple local behaviours achieve a desired global objective, while communicating only with cells within an immediate neighbourhood. 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 (Lovatt et al., 2003, Foreman et al., 2003) or form a spanning tree connecting cells that detected non-critical impacts (Wang et al., 2003, Prokopenko et al., 2004).

Typically, the desirable emergent behaviour on the system level is not easily predictable from local agents' rules and interactions. Our proposed approach to designing localised algorithms is based on an iterative process including the following steps: a) forward simulation leading to emergent behaviour for a class of localised algorithms; b) quantitative measurement of spatiotemporal stability of the emergent patterns; and c) evolutionary modelling of the algorithms, with the metrics obtained at step b) contributing to the fitness functions. This paper investigates the step c).

The next section will cover some background on measures of emergent behaviour, followed by a description, in section 3., of a localised algorithm producing impact boundaries and emergent recovery membranes. Section 4. will then describe a genetic algorithm and our experiments on evolving the recovery membranes to fit the desired criteria.

2. Background

A very promising direction was investigated by Wright et al. (Wright et al., 2000) who designed *a measure of emergence* of swarming/flocking behaviour in multi-agent systems. The approach borrows from the analysis of dynamical systems and assumes a selection of a set of local state variables (eg., 2-dimensional coordinates of agents/particles and their corresponding velocities). The proposed measure Ω estimates the level of self-organisation in the multi-agent flock via approximation of the dynamical system's *characteristic dimension* — i.e., by determining how well a swarm/flock can be described as a single body. The characteristic dimension is approximated through the Shannon entropy of singular value spectra sampled over time (Wright et al., 2000). The proposed calculation involves a number of approximations and is computationally intensive, but importantly, it associates "a systems' ability to exhibit emergent behaviour with sudden transitions in the Ω measure, relative to smooth changes in system parameters" (Wright et al., 2000).

There are other information-theoretic metrics that target complexity of a multi-agent system. For example, a rule-space of 1-dimensional cellular automata was characterised with the Shannon entropy of rules' frequency distribution (Wuensche, 1999). 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 input-entropy variance points to a phase transition as well, indicating the edge of chaos (complexity).

We would like to point out, however, that both the Ω measure and the input-entropy of the CA rule-space rely purely on local state variables that can be associated with each particle/cell, and do not directly capture *inter-agent* connections and interactions. In other words, the Ω measure estimates degrees of freedom in the dynamical system, rather than complexity of spatial inter-connections in a multi-agent network. Similarly, the input-entropy of the CA rule-space traces diversity of

rules used over time and identifies temporally persistent configurations, rather than spatial connectivity among neighboring cells. If one needs to capture emergence and stability of a specific spatial pattern (eg., a continuously connected boundary, a maximally connected sub-graph, a desired geometric or topological shape, etc.), then another type of measure is needed.

It is well-known that graph connectivity can be analysed in terms of the size of the largest connected subgraph (LCS) and its standard deviation obtained across an ensemble of graphs, as suggested by the random graphs theory (Erdös and Renyi, 1961). In particular, 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. Thus, a metric based on LCS variance may capture *spatial* aspects of desired emergent patterns (connectivity), while a metric based on entropy of local agent variables (rules, coordinates, velocities, etc.) may capture *temporal* persistence of emergent behaviour.

Importantly, the approach introduced in the work done by (Wright et al., 2000) suggests a feedback using the Ω measure in *evolving* the desired pattern: swarming behaviour as opposed to fully coordinated "crystalline" behaviour or totally uncoordinated dynamics of independent particles. A feedback to evolvable localised algorithms, based on a quantitative measure of global emergent behaviour, is central to our approach as well. Moreover, the identified distinction between spatial and temporal metrics allows us to implement this feedback systematically, focussing on those evolvable aspects that are important for a specific task (spatial connectivity and/or temporal persistence, size of emergent structures, etc.).

3. Impact Boundaries

This section presents a localised algorithm producing *impact boundaries* in the presence of connectivity disruptions and cell failures resulting from impacts. Typically, the damage on the AAV skin caused by a high energy impact is most severe at the point of impact (an epicentre). It will be assumed that not only the cells at the epicentre are destroyed, but the communication capability of the neighbouring cells may be reduced — e.g., the communication damage may propagate out with an exponential decay to a certain radius. In this case, the damage can be characterised by a probability P_i of an error corrupting a message bit i, dependent on proximity of the affected communication port to the epicentre:

$$P_i = \frac{1}{2}\,(1 - \frac{d}{R})^{\alpha},$$

where d is the distance between the involved communication port and the epicentre of the impact with the radius

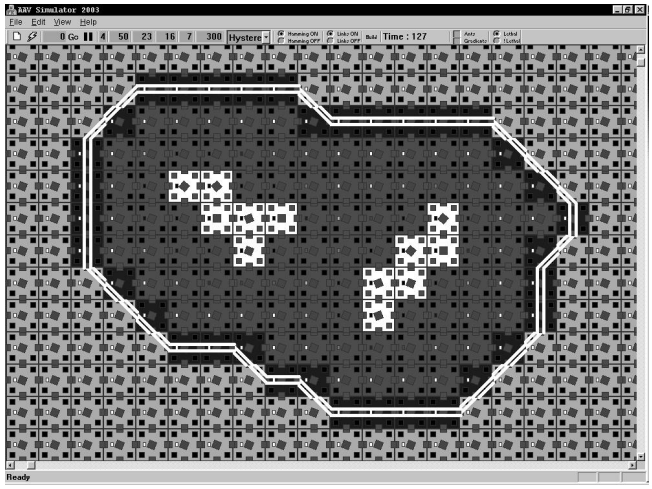

Figure 3: A stable impact boundary: white cells are destroyed, dark-grey cells form "scaffolding", black cells form "frame". Boundary links are shown as white double-lines.

R, and α is the exponential decay of the communication loss. Obviously, multiple impacts result in overlapping damaged regions, and the cumulative bit error probability can be approximated as

$$P_i = \frac{1}{2}\,\sum_{j=1}^{j=m}(1 - \frac{d_j}{R_j})^{\alpha},$$

where m is the number of impacts. The probability that the whole message containing n bits is corrupted depends on the employed error correction code. For example, the $(12, 8)$ Hamming error correction code (8 data and 4 parity check bits) corrects a single bit error per 12 bits, and leads to the following estimation

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

where $r = \frac{1}{12}$ is the error correction rate[1].

The impact boundaries are expected to enclose critically damaged regions and form continuously connected closed circuits (Lovatt et al., 2003). Moreover, impact boundaries must be robust to communication malfunctions caused by proximity to the impact — in other words, in the presence of unreliable communications between cells due to high probabilities $P_{failure}$, and without knowing the distances between each cell and the epicentre. Figure 3 illustrates a typical situation: the white cells are destroyed, normal cells (background colour) do not detect any damage, while cells shown with white double-lines self-organise into the impact boundary. Impact boundaries are intended to ensure a reliable communication pathway around the damaged region, and

[1] A computationally expensive Reed-Solomon error correction code would give a much better error correction rate than $\frac{1}{12}$ in the best case of consecutive "burst" errors.

also create an outline for subsequent self-repair, being inspired by a biological analogy — the clotting of a wound on mammalian skin. Not surprisingly, spatiotemporal stability in impact boundaries is an important concern: even at the periphery of a single impact region the probability $P_{failure}$ may be fairly high. Moreover, multiple impacts often introduce some asymmetry within cell neighbourhoods, so that a cell may be able to transmit but not receive data or vice versa (Lovatt et al., 2003).

3.1 Emergent Membranes and Recovery Mode

In this section, while describing the localised algorithm producing impact boundaries, we shall highlight the role of emergent *recovery membranes* that separate the boundaries from inner impact-surrounding regions, including (possibly recovering) cells that may communicate unreliably. We intend to show that a recovery membrane is an emergent structure, and its emergence is precisely the reason for stabilisation of the impact boundary. In doing so we shall also underscore the parameters used to evolve recovery membranes and stable impact boundaries.

First of all, we consider the following two communication behaviours potentially leading to emergence of an impact boundary, in the absence of sensory readings from cells destroyed by impacts:

(i) At the start of every cycle, every cell sends a *Ping* message to each of its neighbours.

(ii) If a boolean variable *EnableAcks* is true, a cell sends an Acknowledgment reply when it receives a Ping message.

The second behaviour is evolvable, while the first one is pre-determined. Two binary circular arrays are used to store the communication histories for ping messages (*pingArray*) and acknowledgments (*ackArray*) for each communication port. The size of these arrays is called the communication history length.

Every cycle, after receiving communication messages, each cell updates the following parameters:

(1) For each communication port, *PingFailure* will result if the percentage of lost Pings in the *pingArray* is greater than *PingFailThreshold*.

(2) For each communication port, *PingSuccess* will result if the percentage of Pings received in the *pingArray* is greater than or equal to *PingSuccessThreshold*.

(3) For each communication port, *AckSuccess* will result if the percentage of Acknowledgments received in the *AckArray* is greater than or equal to *AckSuccessThreshold*.

The *Failure* and *Success* parameters are hysteretic: they change values only when a sufficient communication history is accumulated. This lagging of an effect behind its cause provides a (temporary) resistance to change and ensures a degree of stability in the treatment of communication connections between any two cells:

(4) A neighbour is considered to be *communicating* when *EnableAcks* is false and *PingSuccess* is true, or both *EnableAcks* is true and *AckSuccess* is true.

(5) *Scaffolding* state S_s will result if there are no *communicating* neighbours.

(6) *Frame* boundary state S_f will result if *PingSuccess* is true for at least one communication port, and *PingFailure* is true for at least one communication port[2].

(7) *Closed* boundary state S_c will result if the cell state is S_f, and there are at least two *communicating* neighbours.

In order for a continuous impact boundary to emerge, the following two communication behaviours were considered:

(iii) if the cell state is S_c, 1) determine a cell α that failed to communicate; 2) determine two communicating neighbour cells β_1 and β_2 nearest to the cell α[3]; 3) map the directions to α, β_1 and β_2 to a direction γ; and 4) send a "Connect(τ, γ)" message to both β_1 and β_2 with a *time to live* parameter τ;

(iv) upon receiving "Connect(τ, γ)" message from a cell α, if the cell state is not S_c, 1) switch to the state S_c, 2) if $\tau > 0$, follow the steps 2) and 3) from (iii) producing new direction γ' and 3) send a "Connect$(\tau - 1, \gamma')$" message to neighbours β_1 and β_2.

The *time to live* τ is an evolvable parameter, and prevents spurious links from persisting. In general, the described policy achieves the desired robustness and continuity of self-organising impact boundaries for a variety of cell shapes (triangular and square) and communication damage probability distributions (Foreman et al., 2003). One particular evolvable behaviour is critical in achieving the desired stability:

(v) If a boolean variable *EnableShutdown* is true, the cell in the *Scaffolding* state S_s will stop transmitting messages.

This is needed in order to break asymmetry between neighboring cells, where a cell is able to transmit data but not receive, or vice versa — it is better in such cases

[2]This condition is similar to the condition of having at least one communicating neighbour and at least one mis-communicating neighbour, but does not assume Acknowledgement messages at all.

[3]When cell δ determines its own neighbour cells nearest to some other neighbour cell α, it chooses the cells β_1 and β_2 on opposite sides (clockwise and counter-clockwise), relative to α.

not to communicate at all. This behaviour is effective in achieving stable impact boundaries. However, the cells that stopped transmitting messages may need to resume communications under certain conditions — for example, when a repair action is initiated, and their neighbours are again 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:

(vi) Recovery state S_r will result if all ports have *pingSuccess* for a number of consecutive cycles, where this number is set by *RecoveryThreshold*.

(vii) Recovery state S_r will result if all ports have *pingFailure* for a number of consecutive cycles, where this number is set by *RecoveryThreshold*.

(viii) A cell stays in the Recovery state S_r and may send communication messages during the next *RecoveryInertia* cycles.

The behaviour (vii) enables totally isolated cells to be in a ready recovery state as well — a feature that, understandably, may be needed when a repair is in progress and cells are being transported across a panel. It is precisely the combination of the "shutdown" and "recovery" conditions that makes the problem of designing (evolving) stable impact boundaries particularly difficult. On one hand, "shutdown" condition (v) disables "asymmetric" messages, making the task of boundary formation easier. On the other hand, conditions (vi)-(viii) may enable "asymmetric" messages, potentially destabilising 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 always forms on the inside of the closed boundary, and on the outside of the recovering area. Interestingly, unlike scaffolding and frame boundary, the membrane is not a designated state into which a cell can switch. Membrane cells shut down their communications, following behaviour (v) like other scaffolding cells, but do not resume communications because behaviours (vi)—(viii) are not applicable, as there are some (but not all) mis-communicating neighbours. Without a membrane, the cells on the frame boundary would be confused by intermittent messages from scaffolding cells attempting recovery. Figure 4 illustrates a checkered-pattern recovery membrane shown with dark-grey colour, while the recovering cells are shown in darker shade of white. It is quite obvious that the membrane cells separate the recovering scaffolding cells from the frame boundary.

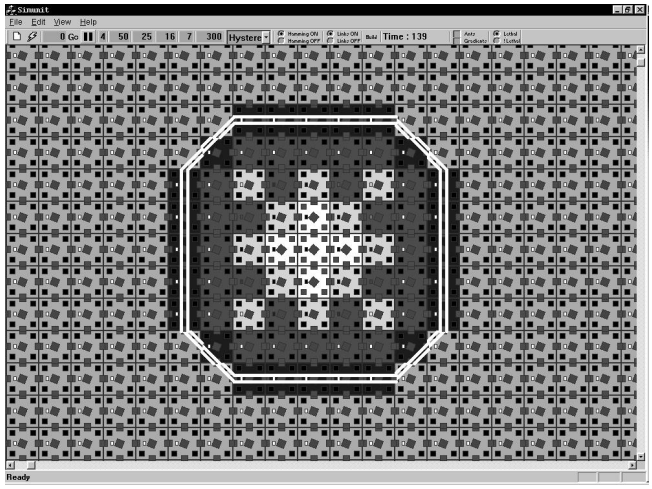

Figure 4: Five white cells at the epicentre are destroyed. A recovery membrane shown in dark-grey "absorbs" scaffolding cells that attempt recovery, shown in darker shade of white, and separates them from the frame, shown in black.

4. Evolving Recovery Membranes

In this section, we describe a Genetic Algorithm (GA) aimed at evolving the parameters used in impact boundary formation. In particular, we examine the interrelationships between communication "shutdown" and "recovery" conditions that give rise to recovery membranes needed for stable impact boundaries. Some of the evolving parameters designate conditionalised branches of the impact boundary algorithm, and some represent various thresholds, so the search space is significantly large. It is well-known that a genetic algorithm, as a problem-solving tool based on biological evolution, works on improving a solution via a search through a process of selection, recombination (crossover) and mutation, and is particularly useful when the search space has many local optima or is too large to use conventional techniques.

4.1 Encoding

We followed a traditional GA encoding: binary strings encoding the chromosome as a structure containing the collection of parameters (genes) and representing a behavioural trait of the individual. Our chromosome contains the following 42 bits:

History, ρ (5 bits) — number of cycles for which a cell remembers the received Pings and Acknowledgements.

PingFailThreshold (7 bits) — a percentage representing the number of Ping messages lost in the last History cycles including the current cycle.

PingSuccessThreshold (7 bits) — a percentage representing the number of Ping messages received in the last *History* cycles including the current cycle.

AckSuccessThreshold (7 bits) — a percentage representing the number of Acknowledgements messages received in the last *History* cycles including the current cycle.

RecoveryThreshold, π (5 bits) — number of cycles the recovery condition must hold for communication recovery to start.

RecoveryInertia (6 bits) — number of cycles in which a cell continues to communicate with its neighbours while recovering.

Time to live, τ (3 bits) — number of times a Connect message is sent before being discarded.

EnableShutdown (1 bit) — a Boolean variable indicating whether a cell shutdown functionality is enabled.

EnableAcks (1 bit) — a Boolean variable indicating whether a cell has the ability to send and receive Acknowledgements messages.

4.2 Fitness/Objective Function

The evolution of recovery membranes is based on spatiotemporal metrics incorporated within a fitness (objective) function. The analysis presented by (Foreman et al., 2003) used two metrics to characterise stability of emergent impact boundaries: spatial and temporal.

The spatial metric is based on the variance in the size of the 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, which is connected to at least one cell in F (an analogue of a maximally connected sub graph or a graph component). We calculate the maximum size $H_{sp}(t)$ of CBF's in self-organising impact boundaries at each cycle. Its variance σ^2_{sp} over time is then used as a spatial metric within the objective function. This metric, as mentioned before, is inspired by random graphs theory and is intended to capture *spatial connectivity* in impact boundaries. A continuous boundary may, however, change its shape over time, without breaking into fragments, while keeping the size of CBF almost constant. Therefore, a temporal metric that measures the diversity of cell transitions (analogously to cellular automata input-entropy or the Ω measure used to characterise flocking behaviour) may be complementary.

In order to analyse *temporal persistence*, we consider state changes in each cell at every time step. Given 6 symmetric boundary links possible in each square cell ("left-right", "top-bottom", "left-top", etc.), there are 2^6 possible boundary states (including "no-boundary"), and $m = 2^{12}$ transitions. The entropy $H_{temp}(t)$ of a particular frequency distribution $S_i(t)$, where t is a time

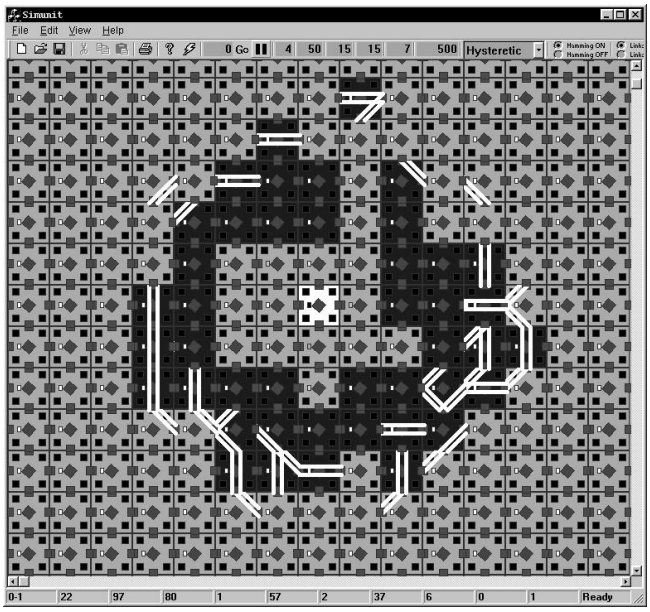

Figure 5: A chaotic boundary with $\overline{H_{sp}} \leq 16$ and zero-length communication ρ. A membrane does not form at all. Both σ^2_{sp} and σ^2_{temp} are low-to-medium.

step, and i is a cell transition index: $1 \leq i \leq m$, can be calculated as follows:

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

where n is the total number of cells, and $S_i(t)$ is the number of times the transition i was used at time t across all cells. Again, the variance σ^2_{temp} of the entropy $H_{temp}(t)$ over time is used as a temporal metric within the objective function.

Our task is complicated by the fact that emergent structures are characterised by a *phase transition* detectable by either σ^2_{sp} or σ^2_{temp}, rather than a particular value range. Therefore, simply rewarding low values for these entropy-based metrics would be insufficient. In particular, it has been observed (Foreman et al., 2003) that both metrics are low-to-medium for algorithms with zero-length communication ρ (tropistic algorithms and chaotic regimes — Figure 5), increase dramatically for ρ in the range $1 \leq \rho \leq \rho_0$, where ρ_0 is a critical value at and below which complex unstable behaviours occur (Figure 6), and undergo a phase transition to very low values when $\rho > \rho_0$ (hysteretic algorithms and ordered regimes).

The critical value ρ_0 is, of course, dependent on all other parameters used by the algorithm. Nevertheless, the chaotic regimes, which are more stable simply due to a small number of connections, can often be identified by a low average $\overline{H_{sp}}$ of the maximum sizes $H_{sp}(t)$ of CBF's in impact boundaries, ruling out at least zero-length histories. In particular, impact boundaries with the average $\overline{H_{sp}} \leq 16$ can be safely ruled out — the

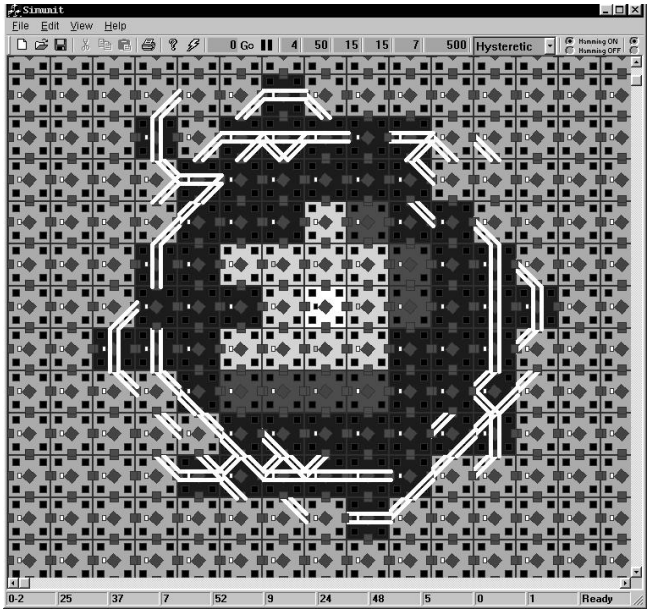

Figure 6: An unstable boundary with ρ close to its critical value. The membrane is fragmentary. Both σ^2_{sp} and σ^2_{temp} are close to their peaks: the phase transition.

resulting chaotic patterns, illustrated in Figure 5, are of no interest.

On the other hand, a preference among ordered regimes towards shorter histories is another useful identifier of a phase transition and the critical value ρ_0. Besides, a shorter communication history ρ enables a quicker response, as do lower values of τ and π. Thus, our first experiment used minimisation of the following objective function:

$$f_{sp}(\beta) = \begin{cases} M & \text{if } \overline{H_{sp}} \leq 16; \\ 4.0\ \sigma^2_{sp} + \rho + \tau + \pi + \beta\ \overline{H_{sp}} & \text{if } \overline{H_{sp}} > 16, \end{cases}$$

where M is the maximal integer value provided by the compiler. The coefficient β reflects the relative importance of the length of impact boundaries in the objective function — sometimes, it may be as important to obtain smallest possible impact perimeter as it is to maintain a shortest possible communication history. We alternated between $\beta_1 = 0.25$ and $\beta_2 = 2.0$.

The second experiment focussed on temporal metric σ^2_{temp} embedded in the objective function $f_{temp}(\beta)$ constructed analogously to the function $f_{sp}(\beta)$:

$$f_{temp}(\beta) = \begin{cases} M & \text{if } \overline{H_{sp}} \leq 16; \\ 10^5\ \sigma^2_{temp} + \rho + \tau + \pi + & \\ \qquad + \beta\ \overline{H_{sp}} & \text{if } \overline{H_{sp}} > 16. \end{cases}$$

Finally, our ultimate objective function is defined as follows:

$$f(\beta) = \begin{cases} M & \text{if } \overline{H_{sp}} \leq 16; \\ \frac{1}{2}\ (4.0\ \sigma^2_{sp} + 10^5\ \sigma^2_{temp}) + & \\ \qquad + \rho + \tau + \pi + \beta\ \overline{H_{sp}} & \text{if } \overline{H_{sp}} > 16 \end{cases}$$

Each experiment involves an impact at a predefined cell, and lasts 500 cycles; the first 30 cycles are excluded from the series $H_{sp}(t)$ and $H_{temp}(t)$ in order not to penalise longer history lengths ρ. We repeat the experiment 3 times for every chromosome and average the objective (fitness) values obtained over these runs.

4.3 Selection

We have chosen a *generation gap* replacement strategy hoping to use better search capabilities offered by the generational replacement and faster convergence typically provided by the steady state selection. Some of this faster convergence, however, may be explained by the stochastic nature of the selection operator: the rate of genetic drift in steady state selection is twice that of generational selection (Rogers and Prügel-Bennett, 1999). In our experiments, we set the generation gap parameter $G = 0.2$. In other words, the entire old population (40 chromosomes) is sorted according to fitness, and we choose the best 20% for direct replication in the next generation, employing an elitist selection mechanism. The selection phase and recombination (crossover) phase can be merged (Thierens and Goldberg, 1994). We follow a similar approach but still explicitly keep the selection phase replicating the elitist offspring. The rest of selection functionality is moved into the crossover.

4.4 Crossover and Mutation

We used a variation of the n-point crossover, where the probability of having n points in the crossover depends on the fitness of the chromosome. We choose this since there are benefits of having a low and high n-point crossover, enabling a better exploration in the search space. Our variation involves three different crossovers, each having an equal probability of contributing to the generation of new chromosomes. In other words, after the elite takes 20% of the new population, for each remaining place we randomly perform a crossover chosen among the following three:

Elitist driven: parent 1 is randomly chosen from the best performing 20% of the old population and parent 2 is randomly chosen from the entire old population, followed by a low 1- to 2-point crossover (the number of points is determined randomly). This low-point crossover is used here since it makes sense to disturb an elitist-driven solution as little as possible.

Mid-range: parent 1 is randomly chosen from the next 20% - 50% of the old population and parent 2 is randomly chosen from the entire old population, followed by a medium 1- to 4-point point crossover (the number of points is determined randomly as well). This medium-point crossover is more applicable when it makes sense to disturb a mediocre solution.

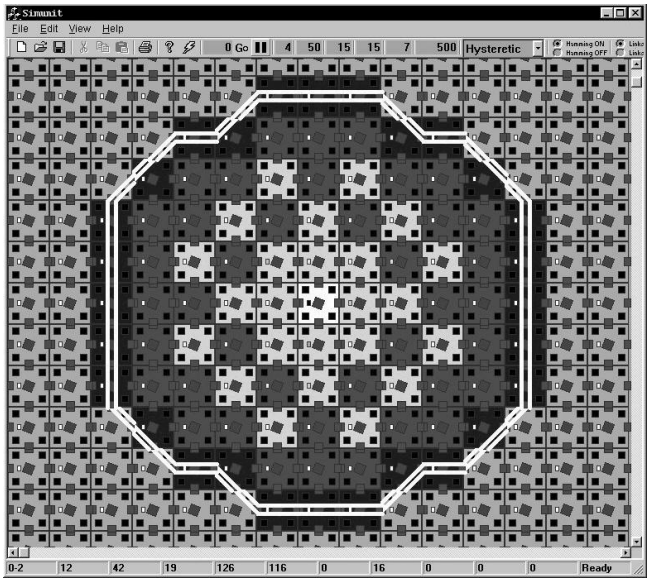

Figure 7: A large checkered-pattern membrane, with short hysteresis, within a morphing but closed and continuous boundary ($\beta = 0.25$). Both σ^2_{sp} and σ^2_{temp} are low.

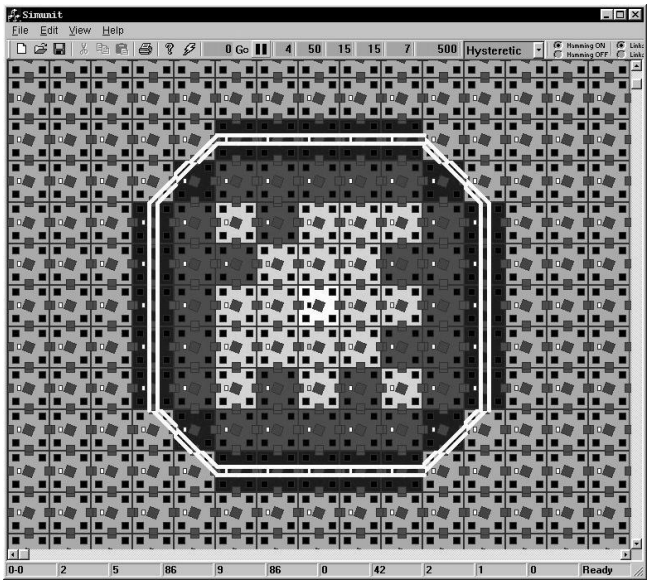

Figure 8: A small membrane, with long hysteresis, within a regular octagonal boundary ($\beta = 2.0$). Both σ^2_{sp} and σ^2_{temp} are very low.

Remainder: parent 1 is randomly chosen from the worst performing 50% of the population and parent 2 is randomly chosen from the entire old population, followed by a high 1- to 7-point crossover, potentially disrupting an under-fit solution a lot in order to enter a new search region.

A feature of this multi-point crossover is that the first gene is always inherited from parent 1. We also ensure that offspring generated from the crossover is unique. We use a slightly higher than typical mutation rate: each bit in the chromosome has a 0.04 probability of being flipped. We also ensure that the mutation results in a unique chromosome by repeating mutation if the produced chromosome already exists in the new population. Mutation is not performed on the chromosome generated via elitist selection.

5. Experimental Results

Our first experiment minimising the objective function $f_{sp}(0.25)$ was mainly concerned with spatially connected and stable impact boundaries, and the form taken by corresponding recovery membranes. The length of the boundary was of lesser importance. Not surprisingly, the evolved solution achieved long robust and continuous impact boundaries with $\overline{H_{sp}} = 40$ (Figure 7), around large impact-surrounding regions, while requiring fairly short hysteresis: $\rho = 2$ and $\pi = 5$. The stabilisation of an impact boundary around a large region occurs at the periphery of the communication damage, where the probability $P_{failure}$ falls to 0 due to the error correction code, and the process has a cascading nature, where the boundary expands to eventually cover all the impact-surrounding region. Interestingly, the evolved continuous boundaries may change their shape and only rarely stabilise as a regular octagon, while keeping their length $\overline{H_{sp}} = 40$ constant. The emergent recovery membrane evolved to separate boundary from the recovering cells. It has a checkered pattern that can be explained by the opposing nature of the conditions (vi) and (vii), and the short hysteresis enabling oscillations between recovery and scaffolding states.

On the other hand, minimisation of $f_{sp}(2.0)$ resulted in more compact impact-surrounding regions ($\overline{H_{sp}} = 32$, Figure 8) and thinner membranes, at the expense of longer hysteresis: $\rho = 6$ and $\pi = 4$. These boundaries morph as well, but generally keep the shape of a regular octagon. Interestingly, the second experiment produced a typical *speciation*, where the longer hysteresis solutions took only one niche, while shorter hysteresis $\rho = 2$ and $\pi = 4$ solutions co-evolved into a separate niche, both niches evolving compact regions with boundaries $\overline{H_{sp}} = 32$. This supports our conjecture that the employed generational gap selection with $G = 0.2$ counteracts the genetic drift to a reasonable degree.

Both solutions favoured $\tau = 1$ as expected for square cells (while triangular cells require at least $\tau = 2$ to achieve continuity). Also, the evolved cells prefer to send Acknowledgements ($EnableAcks = 1$) and stop transmitting in scaffolding state ($EnableShutdown = 1$). Without the latter feature membranes would not emerge, and the recovering cells would disrupt the boundaries.

The second pair of experiments focussed on evolving temporally stable boundaries, minimising $f_{temp}(0.25)$ and $f_{temp}(2.0)$. These experiments produced results very similar to the ones obtained by minimising the spatial

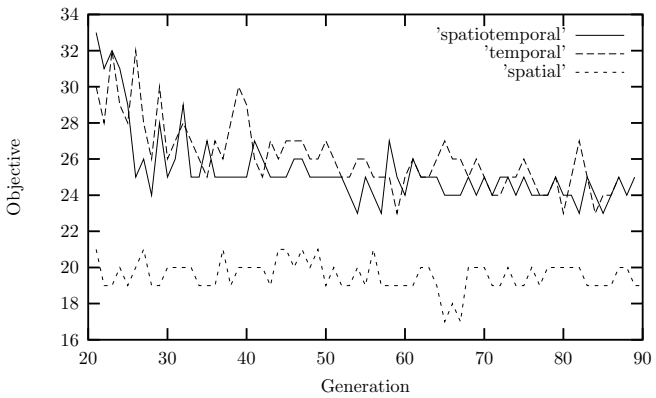

Figure 9: The functions $f_{sp}(0.25)$, $f_{temp}(0.25)$ and $f(0.25)$.

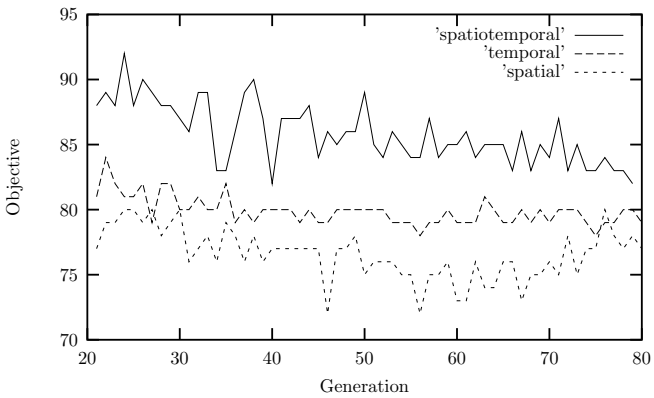

Figure 10: The functions $f_{sp}(2.0)$, $f_{temp}(2.0)$ and $f(2.0)$.

metric. In particular, two types of boundaries evolved, one with a shorter hysteresis ($\rho = 2$ and $\pi = 8$) and longer boundaries $\overline{H_{sp}} = 40$, and the other with a longer hysteresis ($\rho = 6$ and $\pi = 4$) and shorter boundaries $\overline{H_{sp}} = 32$. Both solutions involve membranes, the main difference being that even the shorter hysteresis solutions with longer boundaries tend to reach and retain the regular octagon shape. However, sometimes the boundary fragments. These outcomes are expected for an objective function rewarding temporal stability.

The final pair of experiments combined the spatial and temporal metric, minimising $f(0.25)$ and $f(2.0)$. Again, both types: ($\overline{H_{sp}} = 40$, $\rho = 2$, $\pi = 8$) and ($\overline{H_{sp}} = 32$, $\rho = 6$, $\pi = 8$) were produced. As expected, the evolved membranes and boundaries were more stable, and had a regular octagon shape in both cases, mostly without morphing or fragmenting.

As mentioned above, the case $\beta = 0.25$ results in longer boundaries that are capable of morphing without breaking into fragments. The objective functions $f_{sp}(0.25)$, $f_{temp}(0.25)$ and $f(0.25)$ for the most fit individual in each generation are plotted in Figure 9. All plots exclude the initial period (20 generations) of the rapid decrease typical for GA-based exploration of the search-space. The function $f_{sp}(0.25)$ converges well but does not explore the search-space considerably. This is so because it uses the spatial metric and, therefore, rewards the boundary's continuity, ignoring morphing instability — which is "allowed" in this case. On the contrary, the function $f_{temp}(0.25)$ attempts to minimise morphing instability and has to explore a large part of the space.

The objective functions $f_{sp}(2.0)$, $f_{temp}(2.0)$ and $f(2.0)$ are plotted in Figure 10. The case $\beta = 2.0$ results in shorter boundaries that cannot morph without breaking into fragments, so any instability leads to fragmentation. Consequently, the function $f_{temp}(2.0)$ shows the convergence within a narrower band — it uses the temporal metric and rewards persistence, ignoring occasional fragmentations. Its counterpart, the function $f_{sp}(2.0)$, specifically puts selection pressure on continuity, leading to a wider exploration and poorer convergence.

Importantly, the functions $f(0.25)$ and $f(2.0)$, using spatiotemporal metrics, provided a good compromise in both cases. The experiments confirmed that the choice of an objective function depends on the main task: when the target is not only continuity of impact boundaries but also their shape, then the spatiotemporal metric is more suitable. On the other hand, if morphing is acceptable, then a spatial metric capturing connectivity rather than shape is sufficient. Similarly, a temporal metric may be better suited for a general shape design if sporadic fragmentations are tolerable.

6. Conclusion

In this paper we suggested and verified a methodology underlying design of localised algorithms for complex multi-agent systems, exemplified by self-monitoring aerospace vehicles. We started by briefly describing a multi-agent algorithm leading to emergence of impact boundaries and recovery membranes, followed by quantitative measurement of spatiotemporal stability of the emergent patterns. The graph-theoretic and information-theoretic metrics, capable of identifying phase transitions, contributed to fitness functions for evolutionary modelling of boundaries and membranes. The produced results are promising and demonstrate the possibility for a multi-objective design of localised algorithms. In particular, the desired response time as well as size (and potentially, shape) of impact boundaries and membranes may be specified in advance, leaving the precise logic and parameterisation of the localised algorithms to selection pressures. We believe that the proposed methodology is well suited to the *design at the edge of chaos*, where the design objective (e.g., a specific shape) may be unstable, while other parameters (e.g. the response time) may be optimal. The presented methods should increase the reliability of the design of complex multi-agent systems, accounting for emergent patterns that are not easily predictable by human designers.

Acknowledgements This work was carried 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 for many open and motivating discussions, especially Tony Farmer, Mark Foreman, Mark Hedley, Geoff James, Mark Johnson, Chris Lewis, Geoff Poulton, Don Price, Andrew Scott, Sarath Seneviratne, and Philip Valencia. The authors are also grateful to all members of $GREM^{Lab}$ (a collaborative CSIRO laboratory on Global Response Engineering in Multi-Agent Networks) and the discussion group on "Entropy and self-organisation in multi-agent systems" in CSIRO.

References

Abbott, D., Batten, A., Carpenter, D., Dunlop, J., Edwards, G., Farmer, A., Gaffney, B., Hedley, M., Hoschke, N., Isaacs, P., Johnson, M., Lewis, C., Murdoch, A., Poulton, G., Price, D., Prokopenko, M., Rees, D., Scott, A., Seneviratne, S., Valencia, P., Wang, P., and Whitnall, D. (2003). Development and evaluation of sensor concepts for ageless aerospace vehicles (4): Phase 1 — implementation of the concept demonstrator. In *CSIRO CTIP Report No TIPP 1898*. September 2003.

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., and Winter, J. (2002). Development and evaluation of sensor concepts for ageless aerospace vehicles. development of concepts for an intelligent sensing system. In *NASA technical report NASA/CR-2002-211773*. Langley Research Center, Hampton, Virginia.

Durbeck, L. and Macias, N. (2002). Defect-tolerant, fine-grained parallel testing of a cell matrix. In Schewel, J., James-Roxby, P., Schmit, H., and McHenry, J., (Eds.), *Proceedings of SPIE ITCom 2002 Series, Vol. 4867*.

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

Estrin, D., Govindan, R., Heidemann, J., and Kumar, S. (1999). Next century challenges: Scalable coordination in sensor networks. In *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networks*, pages 263–270. ACM Press.

Foreman, M., Prokopenko, M., and Wang, P. (2003). Phase transitions in self-organising sensor networks. In Banzhaf, W., Christaller, T., Dittrich, P., KIm, J., and Ziegler, J., (Eds.), *Advances in Artificial Life - Proceedings of the 7th European Conference on Artificial Life (ECAL)*, volume 2801 of *Lecture Notes in Artificial Intelligence*. Springer Verlag.

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

Macias, N. and Durbeck, L. (2002). Self-assembling circuits with autonomous fault handling. In Stoica, A., Lohn, J., Katz, R., Keymeulen, D., and Zebulum, R., (Eds.), *Proceedings of NASA/DoD Conference on Evolvable Hardware*. IEEE Computer Society Press.

Prokopenko, M., Wang, P., Foreman, M., Valencia, P., Price, D., and Poulton, G. (2004). On connectivity of reconfigurable impact networks in ageless aerospace vehicles. *Journal of Robotics and Autonomous Systems*, Special Issue:in press.

Rogers, A. and Prügel-Bennett, A. (1999). Genetic drift in genetic algorithm selection schemes. *IEEE Transactions on Evolutionary Computation*, 3(4):298–303.

Thierens, D. and Goldberg, D. (1994). Elitist recombination: an integrated selection recombination ga. In Michalewicz, Z., Schaffer, J., Schwefel, H.-P., Fogel, D., and H., K., (Eds.), *Proceedings of the First IEEE Conference on Evolutionary Computation*, pages 508–512. IEEE Press.

Wang, P., Valencia, P., Prokopenko, M., Price, D., and Poulton, G. (2003). Self-reconfigurable sensor networks in ageless aerospace vehicles. In Nunes, U., de Almeida, A., Bejczy, A., Kosuge, K., and Machado, J., (Eds.), *IEEE Proceedings of the Eleventh International Conference on Advanced Robotics, ICAR 2003*, pages 1098–1103. Coimbra, Portugal.

Wright, W., Smith, R., Danek, M., and Greenway, P. (2000). A measure of emergence in an adapting, multi-agent context. In Meyer, J., Berthoz, A., Floreano, D., Roitblat, H., and Wilson, S., (Eds.), *Proceedings of the Sixth International Conference on the Simulation of Adaptive Behaviour, SAB 2000*, pages 20–27. ISAB Press.

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.