Complexity Metrics for Self-monitoring Impact Sensing Networks

Mikhail Prokopenko[1], Peter Wang[1], Don Price[2]
[1]CSIRO Information and Communication Technology Centre
[2]CSIRO Industrial Physics
CSIRO ICT Centre, Locked bag 17, North Ryde 1670, Australia
e-mail: mikhail.prokopenko@csiro.au, phone: (612)-9325-3264

Abstract

In this paper we describe novel metrics measuring complexity in self-organising networks. The metrics are investigated within the context of decentralised inspections, developed and implemented as part of the joint CSIRO-NASA Ageless Aerospace Vehicle (AAV) research project. The AAV Concept Demonstrator is a hardware multi-cellular sensing and communication network which is expected to detect and react to multiple impacts, without any centralised controllers. We present an extension of an Ant Colony Optimisation algorithm, using an Adaptive Dead Reckoning Scheme and producing robust and reconfigurable minimum spanning trees connecting autonomous AAV cells. We then introduce a new metric detecting emergence through irregularities in the multi-agent communications, and contrast it with conventional macro-level ("global-view") graph-theoretic metrics.

1 Introduction

Structural health monitoring (SHM) is expected to play a critical role in the development and exploitation of future aerospace systems, operating in harsh working environments and responding to various forms of damage and possible manufacturing and/or assembly process variations. NASA's vision of self-monitoring robust aerospace vehicles includes both local and global SHM systems [8]. The local inspections are anticipated to autonomously identify, evaluate, and trigger repair for a wide range of damage or defect conditions in aerospace materials and structures. In parallel, global inspections should be able to dynamically evaluate structural integrity across large and remote areas. This dual architecture, in turn, entails the need for dynamic and decentralised algorithms used in damage detection, evaluation, diagnostics and prognosis.

A distinguishing feature of complex systems, using dynamic, distributed and decentralised algorithms, is the emergence of system-level behaviour out of the interactions among local nodes. Traditional multi-component systems, including SHM systems, do not exhibit self-organisation — instead, they rely on fixed multiple links among the components in order to efficiently control the system, having fairly predictable properties, at the expense of being less scaleable and less robust. Consequently, the traditional design and verification methodologies developed so far have very limited applicability with respect to complex systems: they do not capture self-organisation and cannot fully measure resilience, fault-tolerance and recovery.

A promising new approach to analysis, design and verification of complex SHM systems is to use entropy and information transfer in measuring self-organisation [20, 7, 18, 19, 14]. The metrics for detecting and measuring non-deterministic emergent behaviour can be incorporated within fitness functions, used in evolving desired system-level outcomes [18, 19]. A significant drawback of many available measures, however, is their assumption of a full information of the systems dynamics — a "global view", and a subsequent deterioration of quality when this information is only partially accessible. The agenda of this paper is, therefore, three-fold: we aim to present a dynamic decentralised algorithm solving an SHM task via self-organisation; apply a macro-level ("global-view") metric to capture the quality of the emergent solution; and then verify the solution with a micro-level metric based on the irregularity of the multi-agent communications space. Such a metric will suggest a way to develop metrics based only on partial information — *localisable metrics*.

This investigation is carried out in the context of the CSIRO-NASA Ageless Aerospace Vehicle (AAV) project. The AAV Concept Demonstrator (CD) is a hardware multi-cellular sensing and communication network whose aim is to detect and react to impacts caused by projectiles that, for a vehicle in space, might be micro-meteoroids or space debris [1, 12, 14]. The next section describes the AAV-CD and the problem of decentralised inspections, followed by a review of complexity measures developed over the recent years (Section 3). Section 4 describes a modified Ant Colony Optimisation algorithm pre-optimising decen-

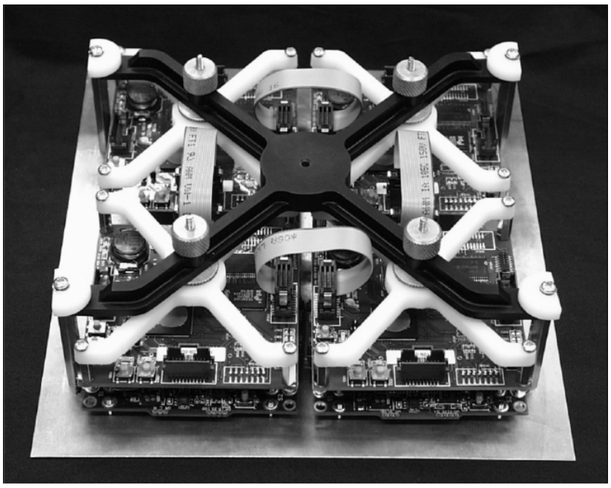

Figure 1. An aluminium panel with 4 cells.

tralised inspections, while Section 5 introduces the employed metrics and presents the experimental results.

2 AAV Concept Demonstrator

The primary principle that is followed in the AAV work is the *emergence of a global response* as a result of interactions involving transfer of local information. This approach attempts to completely avoid or reduce the number of single points-of-failure, leading to *robust* impact sensing networks. The second important principle is *scalability* with respect to manufacturing new or repairing damaged modules or "cells". The AAV-CD consists of cells, that not only form a physical shell for an aerospace vehicle, but also have sensors, logic, and communications. Currently, each cell contains a small number of passive piezoelectric polymer sensors bonded to an aluminium skin panel in order to detect the elastic waves generated in the structure by impacts. At the moment, the structure of the AAV-CD is a hexagonal prism. A modular aluminium frame is covered by 220 mm x 200 mm, 1-mm thick aluminium panels that form the outer skin of the structure. Each such panel contains four cells (Figure 1), and each of the six sides of the prism contains eight of these panels, so the entire AAV-CD contains 48 panels and 192 cells. Each cell also contains 2 digital signal processors, one of which acquires data from the sensors, while the other runs the agent software and controls the communications with its neighbouring cells. Importantly, a cell communicates only with 4 immediate neighbours: the AAV-CD does not employ centralised controllers or communication routers. Single cells may detect impacts and triangulate their locations, while collections of cells may solve more complex tasks; for example, produce an impact boundary with desired characteristics [14] or an impact network [17] to pre-optimise inspections and repairs.

Decentralised inspection across the AAV network array may require an *impact network* — a robust reconfigurable network connecting remote AAV cells that belong to a specific class, e.g., the cells that registered impacts with energies within a certain band (non-critical impacts) [17, 1]. The self-organising impact networks create an adaptive topology allowing inspection agents (communication packets or, potentially, swarming robots) to quickly explore the area and evaluate the damage (e.g., identify densities of impacts typical for a meteor shower) — particularly, where a number of individually non-critical damage sites may collectively lead to a more serious problem. Robotic agents may need an impact network which solves a travelling salesperson problem (TSP). On the other hand, a shortest or *minimum spanning tree* is often required in order to enable decentralised inspections with virtual (software) agents.

Let us define the AAV impact network. The two-dimensional AAV array can be represented by a planar grid graph G: the product of path graphs on m and n vertices $V(G)$, which are points on the planar integer lattice, connected by the edges $E(G)$ at unit distances (Figure 2). Given a number of non-critical impacts, all cells that have detected these non-critical impacts can be represented by a subset P of $V(G)$. We need to identify those edges Z in $E(G)$ which connect the vertices in P minimally, so that the total distance (a sum of unit distances assigned to edges Z) is shortest. This problem is, essentially, the standard minimum spanning tree (MST) problem, except that a spanning tree is defined for a graph, and not for a set of vertices. Our problem is sometimes referred to in literature as the rectilinear minimum (terminal-) spanning tree (RMST) problem, while the vertices in P are called terminals, and is a fundamental problem in VLSI design [9]. The important difference between MST and RMST is that rather than choosing MST edges out of the graph edges $E(G)$ directly connecting pairs of vertices, we need to find multi-edge rectilinear paths between vertices in P, minimising the total distance. In other words, we need to define an auxiliary complete graph A, whose vertex set is P and in which the edge pq for $p, q \in P$ with $p \neq q$ has length equal to the Manhattan distance between nodes p and q. The graph A is not a grid graph — it is an abstraction useful only to formalise an impact network. After a standard MST, A_t, is identified in A, we merely need to convert all edges in A_t to rectilinear paths on the grid graph G (Figure 2).

The impact network problem is complicated by possible "obstacles" created by discontinuities in the AAV grid graph G. Initially, the grid graph G is solid — it does not have any "holes" — so its complement in the infinite orthogonal planar grid is connected. New critical impacts may create such holes in the grid. Figure 2 illustrates the RMST problem with two scenarios. The first case is shown in the top part, and involves three edges and a simple MST, A_t. The second case is shown in the bottom part: some cells are destroyed (the corresponding vertices are removed), and

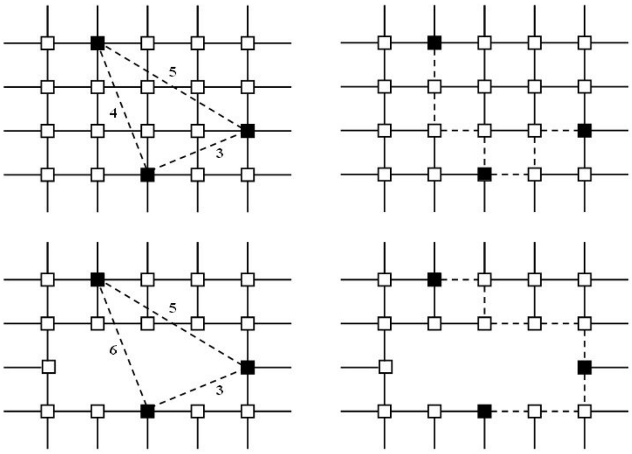

Figure 2. Three impact nodes are shown in black. The top-left figure shows a complete auxiliary graph A (dashed lines) with 3 edges. The conversion of its MST to rectilinear paths is shown in the top-right figure (dashed edges). Two lower figures show the graph with some vertices removed: the bottom-left figure shows an updated auxiliary graph A, and the bottom-right figure shows conversion of the new MST to rectilinear paths.

the auxiliary complete graph should be updated because one shortest path has changed. This requires a re-computation of its MST, with another edge being selected and converted to a rectilinear path. Therefore, a new obstacle may not just require that a new shortest path is found between the two involved cells (the problem investigated by Wu et al. [21]), but rather that the whole MST is re-evaluated. Moreover, there are cases when a terminal is no longer needed to be included in the RMST, or a new terminal needs to be added.

Thus, from a graph-theoretic standpoint, the representation of the impact network problem changes over time due to insertion of new nodes (e.g., non-critical impacts) or deletion of old nodes no longer fitting the impact range, while the problem's properties change due to varying connection costs (e.g., critical impacts destroying existing paths). In short, we need a dynamic and decentralised computation of a rectilinear minimum terminal-spanning tree in the presence of obstacles. If the information (such as the auxiliary graph A) was available in one central point, then RMST problem would essentially become MST problem, with a subsequent conversion to rectilinear paths. In our case, the auxiliary graph A is not even known at any single node/cell. So, on the one hand, the desired algorithm should be both decentralised and fully dynamic (both insertions and deletions must be handled online), while on the other hand, our main concern is the communication overhead of determining the desired topology.

Our ultimate goal is to develop and verify the methodology for evolving such algorithms — in other words, we should investigate not only algorithmic details, but also evaluation metrics identifying distinct phases in algorithm performance. These metrics can serve as fitness functions, identifying critical parameters [18] and guiding the design towards desired objectives.

3 Background and Motivation

Self-organisation is typically defined as the evolution of a system into an organised form in the absence of external pressures. For example, a self-assembly of network-like structures connecting a set of nodes without using pre-existing positional information or long-range attraction of the nodes is described by Schweitzer and Tilch [15] using Brownian agents that are capable of producing different local chemical information and responding to it in a non-linear manner. These agents solve two tasks *in parallel*: (i) the detection of the appropriate nodes, and (ii) the establishment of stable links between them. A concrete and detailed implementation of fault-tolerant circuit synthesis on a self-configurable hardware platform is provided by the Cell Matrix approach [5]. The approach employs local nodes (Supercells), performing fault detection, fault isolation, configuration of new Supercells, determination of inter-cell wiring paths, and implementation of the final desired target circuit.

In general, self-organising solutions depend on selection pressures which, through their contribution to the evolutionary fitness functions, constrain the emergent behaviour. One example of a generic selection pressure is the *spatiotemporal stability* of emergent patterns: arguably, any pattern has to be stable before exhibiting another useful task-oriented feature. The use of spatiotemporal stability in evolving AAV impact boundaries (continuously connected multi-cellular circuits, self-organising in presence of cell failures and connectivity disruptions around damaged areas) is described in [7, 14, 18], employing information-theoretic and graph-theoretic measures in separating chaotic regimes from ordered dynamics.

The information-theoretic *temporal* AAV metric captured the diversity of rules invoked by AAV cells during an impact boundary formation, and was modelled on the classification of a cellular automata (CA) rule-space, characterised with the Shannon entropy of the rules' frequency distribution [20]. The input-entropy settles to fairly low values for ordered dynamics, but fluctuates irregularly within a narrow high band for chaotic dynamics. For complex CA, 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. Complex dynamics exhibits high variance of input-entropy, pointing to a phase transition between chaos and order.

The graph-theoretic *spatial* AAV metric captured the impact boundary's connectivity in terms of the size of the average connected boundary fragment — an analogue of a largest connected sub-graph (LCS) and its standard deviation over time. The parameter driving the dynamics was a length of communication history μ allowed to be kept by each AAV cell. According to the random graphs theory [6], critical changes are expected to 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. Similarly, an impact boundary is highly disconnected (chaotic phase) when the parameter μ is below its critical value μ_0, varies in connectivity widely (the "edge of chaos") when the parameter μ reaches its critical value μ_0, and becomes well-connected (the ordered phase) when $\mu > \mu_0$. In other words, only complex dynamics exhibits high variance, and the peak of this variance points to a phase transition [14].

Network *connectivity* is, then, another example of an independent selection force rewarding specific multi-agent network topologies. This force, we believe, is related to both efficiency and robustness, which were identified in [16] as critical measures underlying optimal network structures. A very promising direction was investigated by Wright et al. [19] who designed a measure of emergence of multi-agent swarming/flocking behaviour as opposed to both fully coordinated "crystalline" behaviour and totally uncoordinated dynamics of independent particles. The proposed measure Ω estimates the level of self-organisation 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.

Typically, the metrics targeting temporal stability and spatial connectivity require a global view: full information on either cells' states (to determine their diversity) or their inter-connections (to determine sub-graph connectivity). Our specific goal, however, is a metric that can work with partial information, obtained locally. For example, a single temporally stable or a spatially connected boundary fragment is clearly insufficient to claim that the whole structure is stable or connected. We intend to show that the first step towards a localisable metric can be provided by a metric operating not within AAV multi-cellular rule-space or on the surface of the AAV graph, but rather within the inter-agent communication space. Before describing the new metric, however, we need to review some details of the localised algorithm that produces the impact networks.

4 Adaptive Impact Networks

The RMST problem on the AAV skin changes concurrently with the problem-solving process, suggesting that it can be efficiently tackled by Ant Colony Optimisation (ACO) algorithms, proposed and enhanced over recent years by Dorigo and his colleagues [2, 3, 4], 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 [4]. 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* behaviour — *allelomimesis*: the probability with which an ant chooses a trail increases with the number of ants that chose the same path in the past. The process is thus characterised by a positive feedback loop [3].

In the AAV-CD the ants are implemented as communication packets, 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 shutdown) communication links, so critically-impacted cells form obstacles, and the ants are supposed to find the shortest paths around them using positively reinforced pheromone trails. For our problem it is impractical to use two types of pheromone (e.g., the "nest" and "food" pheromones) because each impact cell (node) serves both as a "nest" and a "food" source. Therefore, having two types of pheromone per node would have created multiple pheromone fields, combinatorially complicating the network. In addition, dissipation of pheromone over large distances is not practical either, as it would lead to "flooding" of the network with messages. Hence, the algorithms developed for the AAV network use only one type of non-dissipative evaporating pheromone.

The algorithm presented in [17, 1] 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 ant-like agents. Following [14], we summarise here a main variant of this algorithm, without an IGF, and relying only on DRS. The behaviour of exploring ants includes the following:

(E1) *each impact node generates $K = 4$ exploring ants every $T = 20$ cycles; each ant has a "time to live" counter $\tau_k = 255$, 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 non-zero trail intensity;*

(E3) *if an exploring ant can move to a cell with a non-zero trail intensity, the destination cell is selected according to transitional probabilities [14];*

(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 parameters K, T, τ_k may vary, and the frequency K/T, in particular, has an effect on convergence and the communication overhead [1]. The DRS requires that each 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 tabu 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, using the following rules:

(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 coordinates(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 (minimisation 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 [14]; upon this selection the ant keeps to the chosen path until the obstacle is avoided, as recognised by comparison of current y/x ratio with g;*

(R3) *each cycle, a returning ant deposits pheromone in the quantity inversely proportional to the traversed return distance q (q is incremented by 1 each cycle); the deposited pheromone is limited by a pre-defined maximum φ_{max}.*

The pheromone is deposited on the cells themselves rather than communication links — we deal with pheromone trail intensities φ_j at the cell j. Each cell stores an ant-routing table, used in determining which neighbour cell should be chosen by an incoming ant packet to continue their travel. At any given time point t, the ant-routing table $a_{i,j}(t)$ of node i with respect to all its neighbor nodes j, contained in the neighbourhood N_i, is proportional to pheromone intensities $\varphi_j(t)$. The intensity of trail $\varphi_j(t)$ on the node j gives information on how many ants have traversed the node in the past, and is updated each time an ant agent k passes through the node:

$$\varphi_j(t) = min(\varphi_j(t) + \frac{\sigma_k}{q_k(t)},\ \varphi_{max}),$$

where σ_k is a constant quantity specified for each generated ant k, q_k is the distance traversed by the ant k, and φ_{max} is a limit on pheromone trail intensity. Intuitively, the quantity σ_k represents a pheromone reserve of the ant k, consumed during the return trip (not unlike the well-known ANT-quantity scheme). At the beginning of each cycle, the pheromone evaporates at the rate $\rho \in (0, 1)$:

$$\varphi_j(t) = (1 - \rho)\ \varphi_j(t) = \psi\ \varphi_j(t)\ ,$$

where $\psi = 1 - \rho$ is the pheromone retention rate.

An improvement to the DRS algorithm included an *adaptive* pheromone reserve quantity σ_k and a "time to live" counter τ_k. The pheromone reserve is adaptively allocated by the generating node, based on the ants returned to the node in the past: $\sigma_k = max(\gamma_1\ \hat{q},\ \sigma_{min})$, where $\hat{q}$ is the minimal distance traversed by the returned ants, γ_1 is a scaling factor, and σ_{min} is a lower limit for the pheromone reserve allocated for an ant. Analogously, $\tau_k = min(\gamma_2\ \hat{q},\ \tau_{max})$, where τ_{max} is an upper limit for the counter, and γ_2 is a scaling factor. The Adaptive Dead Reckoning Scheme (ADRS) contributes to a faster *reconfiguration* of trails and minimum spanning trees.

The DRS algorithm produces minimum spanning trees, resulting in reconfigurable impact networks, and performs well in dealing with two well-known problems: the blocking problem and the shortcut problem. *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. The *shortcut* corresponds to a new shorter trail becoming available due to repaired cells. In this section we present a local heuristic (the "pause" heuristic), which contributes to a better convergence of the DRS algorithm. Let us consider probabilistic decisions of a returning ant in the situation when an obstacle blocks a DRS path towards the home node. The DRS algorithm uses only one type of pheromone (the "impact" pheromone). For example, a returning ant facing an obstacle ahead and excluding a backtrack possibility has a $50/50$ chance of turning left or right, when the trails are not yet established. Choosing a direction at this decision node results in the ant depositing the pheromone either on the left or the right node. Clearly, this deposit is not an informed choice, being driven by a $50/50$ chance, and may in fact obscure the pheromone trail. The update of the pheromone on both left and right nodes should, in fact, be done *only* by the ants going in the opposite direction, as these ants have traversed an alternative path. This dilemma is not present when the ants use two types of pheromones. A simple solution enhancing the DRS algorithm, using only one pheromone type, is provided by the "pause" heuristic:

(R4) *an ant, facing an obstacle at cycle t and making a transition to the next node, does not deposit any pheromone at cycle $t + 1$, resuming deposits only from cycle $t + 2$.*

The "pause" heuristic initially produces gaps in the trails, next to each decision point (Figure 3). However, these gaps are eventually filled by the ants going in the opposite direction, leading to the reinforcement of the shortest trail. Figures 3 and 4 illustrate this dynamics with snapshots of the simulated AAV-CD network array.

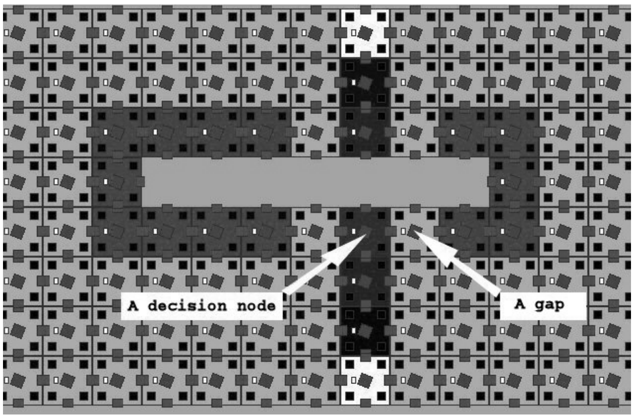

Figure 3. White cells detected non-critical impacts. An initial vertical trail is destroyed by a horizontal obstacle (seven cells are removed). The returning ants explore two alternative possibilities. The gaps in both trails form next to each decision node.

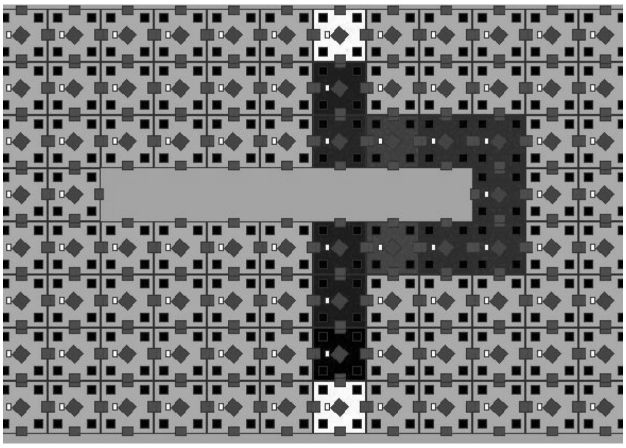

Figure 4. A shorter trail is established.

5 Complexity Metrics and Experiments

The "global-view" or macro-level analysis of ACO-ADRS algorithm convergence is based on the concept of a connected trail-fragment (CTF) — an obvious analogue of a connected sub-graph. A CTF is a set F of cells with $\varphi \geq \vartheta$ (where ϑ 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 which is connected to at least one cell in F. Tracing the average size S of CTFs and its variance σ^2 over time allowed us to identify emergence of the shortest path as a phase transition in network connectivity.

We carried out 10 experiments for each value of the pheromone retention rate ψ in the range between 0.81 and 0.99. During each experiment, a simple straight trail was initially formed between 2 impacts, and then broken at cycle 200 by an obstacle. As before, we calculated the average size of CTF's in impact networks, $S(\psi)$, at each time-

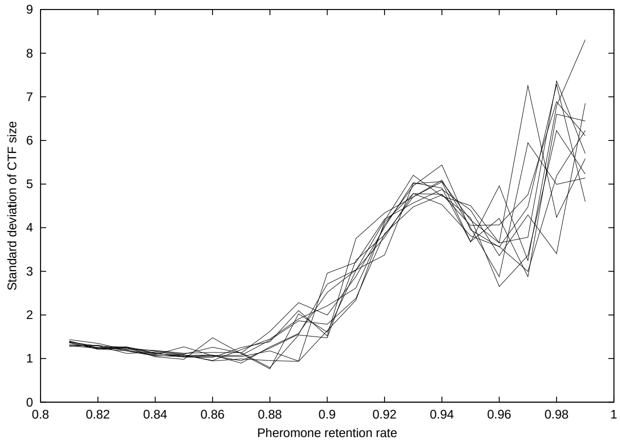

Figure 5. A chaotic phase (fragmentary trails) is separated by the edge of chaos (first maximum at $\psi = 0.94$) from the first ordered phase (stable short trails at $\psi \approx 0.96$); with another phase transition to combined trails ($\psi > 0.98$).

point, and its standard deviation, $\sigma(\psi)$, over time. The expected three types of dynamics: chaotic, complex, and ordered were observed. Two "ordered" phases are observed (Figure 5): the first (and the one we are interested in) is the emergence of the stable shorter trail around the obstacle as opposed to the longer trail, followed by the emergence of both stable trails around the obstacle (with a higher combined length). The first "ordered" phase is separated from the chaotic phase, $\psi < 0.94$, by the "edge of chaos", $\psi \in [0.94, 0.96]$, and is identified by the minimum of $\sigma(\psi)$, at the retention rate $\psi \approx 0.96$. There is another ("edge of order") region of complexity, $\psi \approx 0.97$, preceding the second "ordered" phase at very high retention rates $\psi > 0.98$. This phase is of no interest: at such rates there is enough pheromone to support many trails.

Thus, the spatial metric $\sigma(\psi)$ suggests that, in terms of solving the blocking problem, the optimal pheromone retention rate ψ can be identified as the one which attains the minimum of the standard deviation $\sigma(\psi)$, following the "edge of chaos" pointed to by the first maximum of $\sigma(\psi)$, as we increase ψ. In other words, well-connected impact networks emerge in the ordered phase, characterised by minimal $\sigma(\psi)$ and lower entropy of the network.

Self-organisation may seem to contradict the second law of thermodynamics that captures the tendency of systems to disorder. The "paradox" has been explained in terms of multiple coupled levels of dynamic activity (the Kugler-Turvey model [10]) — self-organisation and the loss of entropy occurs at the macro-level, while the system dynamics on the micro-level generates increasing disorder. One convincing example is described by Parunak and Brueckner [11] in context of pheromone-based coordination. Their work defines a way to measure entropy at the macro level (agents' behaviours lead to orderly spatiotemporal patterns) and micro level (chaotic diffusion of pheromone molecules). In other words, the micro level serves as an entropy "sink" — it permits the overall system entropy to increase, while allowing self-organisation to emerge and manifest itself as coordinated multi-agent activity on the macro level. Another example relates a macro-level increase of coordination potential within a multi-agent team, indicated by a macro-level decrease in epistemic entropy of agents' joint beliefs, with a micro-level increase in the entropy of the multi-agent communication space [13]. Similarly, we intend to show that the emergence of well-connected impact networks, indicated by the minimal variance of their fragments (an approximation of the network heterogeneity), is explained by increased irregularity on a micro-level. This micro-level is the communication space where the inter-agent messages are exchanged.

A characterisation of the micro-level (the entropy "sink") can be obtained if one estimates the "regularity" of the communication space. Let $\alpha_i(t)$ denote the number of ant packets received by the cell i at time cycle t. Then the average load carried by the cell i is given by $\overline{\alpha_i} = \sum_{t=1}^{\Omega} \alpha_i(t)/\Omega$, where Ω is the total number of cycles. The regularity of the series $\alpha_i(t)$ can be measured with the auto-correlation function of an integer delay τ:

$$\gamma_i(\tau) = \frac{\sum_{t=\tau+1}^{\Omega}[\alpha_i(t-\tau) - \overline{\alpha_i}]\,[\alpha_i(t) - \overline{\alpha_i}]}{\sum_{t=1}^{\Omega}[\alpha_i(t) - \overline{\alpha_i}]^2} .$$

The auto-correlation function is equivalent to the power spectrum in terms of identifying regular patterns — a near-zero auto-correlation across a range of delays would indicate high irregularity, while auto-correlation with values close to one indicate very high regularity. The following inverse average is a good approximation of the total irregularity, or *volume-per-channel complexity*, on the micro-level (the communication space):

$$\lambda(\tau) = \frac{M}{\sum_{i=1}^{M} \gamma_i(\tau)} ,$$

where M is the number of cells. It is important to realise that the volume-per-channel complexity metric $\lambda(\tau)$ is a more refined measure than a similar statistic $\xi(\tau) = 1/\gamma_\beta(\tau)$, for a joint series $\beta(t)$ of all ant packets received by all M cells at time cycle t, where $\gamma_\beta(\tau)$ is the auto-correlation function for the joint series $\beta(t)$. The difference between $\lambda(\tau)$ and $\xi(\tau)$ is that the former estimates the regularity of the communication channel employed by each cell $\gamma_i(\tau)$ and then inverts the average over all channels, while the latter is defined in terms of the regularity of the entire communication space $\gamma_\beta(\tau)$. Our conjecture is that the highest irregularity on the micro-level, $\lambda(\tau, \psi)$ over the range of possible values of τ, corresponds to the lowest variance on the macro-level, the minimal $\sigma(\psi)$, and

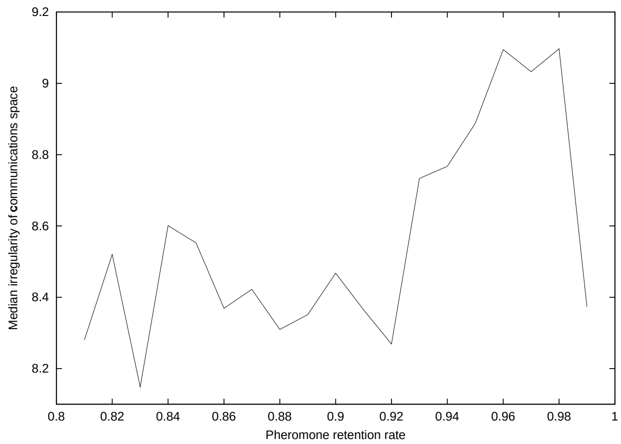

Figure 6. The median $\lambda(20, \psi)$ metric. The maximum at $\psi \approx 0.96$ identifies the most irregular communication space, pointing to emergence of well-connected impact networks.

indicates the observed distinct phases as well. The plot of the median $\lambda(20, \psi)$ calculated during the experiments for each value of the pheromone retention rate $\psi \in [0.81, 0.99]$ is shown in Figure 6. As expected, the maximum is attained at $\psi \approx 0.96$, pointing to the most irregular communication space and indicating a phase transition in the communication space. This retention rate is the rate at which impact networks become well-connected, also identified by the spatial metric σ as the ordered phase. Furthermore, the point $\psi \approx 0.98$ indicates the beginning of the move towards the second ordered phase in the parameter-space where the impact networks include all possible trails. Informally, the most "irregular" communication space, measured by the micro-level metric λ, corresponds to the most "coordinated", well-connected, phase of impact networks. A potential advantage of the λ metric over the spatial metric σ is, however, its possibility to become localised. It may be possible to meaningfully calculate λ over a subset of communication channels, while measuring partial connectivity of an impact network is less likely to succeed.

6 Conclusions

In this paper, we considered emergence of an impact network pre-optimising decentralised inspections on an AAV skin, using an ACO algorithm enhanced with the adaptive dead reckoning and pause heuristics. The modified ACO-ADRS algorithm is deployed in the AAV-CD and robustly solves blocking and shortcut problems, producing rectilinear minimum spanning trees for impact sensing networks. This algorithm exemplifies dynamic decentralised algorithms solving SHM tasks via self-organisation, and we applied two metrics to evaluate the emergent solutions. The spatial metric, $\sigma(\psi)$, measures the quality of impact networks on the macro-level through the connectivity of resultant spanning trees. The volume-per-channel complexity metric, $\lambda(\psi)$, verifies the solution on the micro-level (the multi-agent communications space), and suggests a way to develop metrics based only on partial information — *localisable metrics*.

While we have not evolved parameters for the ACO-ADRS algorithm, the observed phase transitions clearly identify the critical values that would be chosen as a result of selective pressures (spatial stability and/or complexity of communication space) — for example, by a genetic algorithm rewarding stable pheromone trails or irregular communication patterns. This is a subject of future work.

Acknowledgements A part of this work was carried out under CSIRO contract to NASA. It is a pleasure to record our appreciation to Dr. Edward Generazio and Dr. William Prosser (NASA Langley Research Center) for their encouragement and support. The authors are grateful to other members of the AAV project, especially Tony Farmer, Mark Foreman, Vadim Gerasimov, Mark Hedley, Nigel Hoschke, Mark Johnson, Chris Lewis, Geoff Poulton, Andrew Scott and Philip Valencia.

References

[1] Abbott, D., A. Batten, D. C. Carpenter, J. B. Dunlop, G. C. Edwards, A. J. Farmer, B. Gaffney, M. Hedley, N. Hoschke, P. Isaacs, M. E. Johnson, C. J. Lewis, A. Murdoch, G. T. Poulton, D. C. Price, M. Prokopenko, D. E. Rees, D. A. Scott, S. Seneviratne, P. Valencia, P. Wang, and D. F. Whitnall. Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles. Report 4: Phase 1. Implementation of the Concept Demonstrator. CSIRO Telecommunications and Industrial Physics. Confidential Report No. TIPP 1898. September 2003.

[2] Colorni, A., M. Dorigo, and V. Maniezzo. Distributed optimization by ant colonies. In Proc. of the First European Conf. on Artificial Life, Paris, 1991.

[3] Dorigo, M., V. Maniezzo, and A. Colorni. The Ant System: Optimization by a colony of cooperating agents *IEEE Transactions on Systems, Man, and Cybernetics*, B, 26, 1, 1996.

[4] Dorigo, M. and G. Di Caro. Ant Colony Optimization: A New Meta-Heuristic. In Proceedings of 1999 Congress on Evolutionary Computation, Washington DC, July, 1999.

[5] Durbeck, L.J.K., and N.J. Macias. Defect-tolerant, fine-grained parallel testing of a Cell Matrix. In Proc. of SPIE ITCom 2002 Series, Vol. 4867, 2002.

[6] Erdos, P., and A. Renyi. On the strength of connectedness of random graphs. Acta Mathematica Scientia Hungary 12, 261–267, 1961.

[7] Foreman, M., M. Prokopenko, P. Wang. Phase Transitions in Self-organising Sensor Networks. In Proceedings of the Seventh European Conference on Artificial Life, 781–791, Dortmund, September 2003.

[8] E. Generazio. Nondestructive Evaluation Science in the year 2015+. NASA Langley Research Center. Nondestructive Evaluation Sciences Branch. http://nesb.larc.nasa.gov/nde2015.html

[9] Kahng, A.B. and G. Robins. *On Optimal Interconnections for VLSI.* Kluwer Academic Publishers, Boston, MA, 1995.

[10] Kugler, P.N., and M.T. Turvey. *Information, Natural Law, and the Self-Assembly of Rhythmic Movement.* Lawrence Erlbaum, 1987.

[11] Parunak, H.V.D., and S. Brueckner. Entropy and self-organization in multi-agent systems. In Proceedings of the 5th International Conference on Autonomous agents, Montreal, 2001.

[12] Price, D., A. Scott, G. Edwards, A. Batten, T. Farmer, M. Hedley, M. Johnson, C. Lewis, G. Poulton, M. Prokopenko, P. Valencia, and P. Wang. An Integrated Health Monitoring System for an Ageless Aerospace Vehicle. In Proceedings of the 4th International Workshop on Structural Health Monitoring, 310–318, Stanford, September 2003.

[13] Prokopenko, M., and P. Wang. Evaluating Team Performance at the Edge of Chaos. In Polani, D., Browning, B., Bonarini, A., and Yoshida, K. (eds). *RoboCup 2003: Robot Soccer World Cup VII*, LNCS 3020, 89–101, Springer, 2004.

[14] Prokopenko, M., P. Wang, D. Price, P. Valencia, M. Foreman, T. Farmer. Self-organising Hierarchies in Sensor and Communication Networks. To appear in the Artificial Life Journal, Special issue on Dynamic Hierarchies, 2005.

[15] Schweitzer F., and B. Tilch. Self-Assembling of Networks in an Agent-Based Model. *Physical Review E*, 66, 026113 (1-9), 2002.

[16] Venkatasubramanian, V., S. Katare, P.R. Patkar, and F.-P. Mu. Spontaneous Emergence of Complex Optimal Networks through Evolutionary Adaptation. *Computers and Chemical Engineering*, vol. 28 , no 9, 1789–1798, 2004.

[17] Wang, P., P. Valencia, M. Prokopenko, D. Price, and G. Poulton. Self-reconfigurable sensor networks in ageless aerospace vehicles. In Proceedings of the 11th International Conference on Advanced Robotics, 1098–1103, Coimbra, Portugal, 2003.

[18] Wang, P., and M. Prokopenko. Evolvable Recovery Membranes in Self-monitoring Aerospace Vehicles. In Proceedings of the 8th International Conference on the Simulation of Adaptive Behaviour, 509–518, Los Angeles, USA, 2004.

[19] Wright, W.A., R.E. Smith, M. Danek, and P. Greenway. A Measure of Emergence in an Adapting, Multi-Agent Context. In Proceedings of the 6th International Conference on the Simulation of Adaptive Behaviour, 20–27, 2000.

[20] Wuensche, A. Classifying Cellular Automata Automatically: Finding gliders, filtering, and relating space-time patterns, attractor basins, and the Z parameter. *Complexity*, Vol. 4, 3, 47–66, 1999.

[21] Wu, Y.-F., P. Widmayer, M.D.F. Schlag, C.K. Wong. Rectilinear Shortest Paths and Minimum Spanning Trees in the Presence of Rectilinear Obstacles. IEEE Transactions on Computers, vol. C-36, 321–331, 1987.