NASA/CR-2010-216186

Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles

Report 6 – Development and Demonstration of a Self-Organizing Diagnostic System for Structural Health Monitoring

Adam Batten, Graeme Edwards, Vadim Gerasimov, Nigel Hoschke, Peter Isaacs, Chris Lewis, Richard Moore, Florien Oppoizer, Don Price, Mikhail Prokopenko, Andrew Scott, Peter Wang
CSIRO Industrial Physics, Lindfield, New South Wales, Australia

January 2010

NASA STI Program . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NASA Aeronautics and Space Database and its public interface, the NASA Technical Report Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include creating custom thesauri, building customized databases, and organizing and publishing research results.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at *http://www.sti.nasa.gov*

- E-mail your question via the Internet to help@sti.nasa.gov

- Fax your question to the NASA STI Help Desk at 443-757-5803

- Phone the NASA STI Help Desk at 443-757-5802

- Write to:
  NASA STI Help Desk
  NASA Center for AeroSpace Information
  7115 Standard Drive
  Hanover, MD 21076-1320

NASA/CR-2010-216186

Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles

Report 6 – Development and Demonstration of a Self-Organizing Diagnostic System for Structural Health Monitoring

Adam Batten, Graeme Edwards, Vadim Gerasimov, Nigel Hoschke, Peter Isaacs, Chris Lewis, Richard Moore, Florien Oppoizer, Don Price, Mikhail Prokopenko, Andrew Scott, Peter Wang
CSIRO Industrial Physics, Lindfield, New South Wales, Australia

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

Prepared for Langley Research Center
under Purchase Order NNL05AI81P

January 2010

Available from:

NASA Center for AeroSpace Information (CASI)
7115 Standard Drive
Hanover, MD 21076-1320
443-757-5802

Table of Contents

1. Introduction

CSIRO, with support from NASA, has been developing an experimental structural health monitoring (SHM) concept demonstrator and test-bed system (CD) for the detection of high-velocity impacts, such as may occur due to the impact of a micrometeoroid on a space vehicle. The distinguishing feature of this system is that its architecture is based on a complex multi-agent system and its behaviours and responses are developed through self-organisation. It has no central controller. This approach endows the system with a high degree of robustness, adaptability and scalability. Appendix A contains a list of reports and publications that have arisen from this work, and which outline the background to it.

The test-bed has been built as a tool for research into sensor design, sensing strategies, communication protocols, and distributed processing using multi-agent systems. High-velocity impacts are simulated using short laser pulses and/or steel spheres fired using a light-gas gun. Repeatable low-velocity impacts are produced using a pendulum. The SHM system has been designed to be highly flexible: by replacing the sensors and their associated interface and data acquisition electronics, the system can be readily reconfigured for other applications.

The most recent formal report to NASA [1], in April 2004, described the development of the hardware and software of the basic CD system of embedded piezoelectric sensors and processing electronics distributed throughout the structure, along with some multi-agent algorithms to characterise impacts and subsequent damage. Further development of the basic system has since been carried out, but the main focus of the work reported here is the development of a robotic mobile agent, which can move over the outer surface of the CD skin as an independent agent of the self-organising system. This mobile agent can carry out sensing functions that the embedded agents cannot, and it is able to communicate with embedded agents of the CD structure in its vicinity. It is envisaged to be the forerunner of, eventually, a swarm of small robots that will cooperatively perform both inspection and repair functions.

Statement of Work and response

This report outlines work done under NASA Langley Research Center contract no. NNL05AI81P of 23rd September 2005. The Statement of Work requires the following.

"The aim of this project is to develop and critically examine concepts for integrated smart sensing and communication systems that could form the distributed sensing function of a smart vehicle. Such an integrated system may include components deployed for structural monitoring on an NDE agent.

Specific Requirements

Development of the sensing and intelligent decision-making capabilities of the Concept Demonstrator. The contractor shall:

- improve damage evaluation with active ultrasonics so that the impact energy is identified by the concept demonstrator.
- develop and implement a secondary communications modality (wireless) so the concept demonstrator is not solely dependant on wire communications.

- integrate a single mobile (robotic) agent into the concept demonstrator sensing and communication network, to a stage where the agent's movement is determined cooperatively with the fixed cells.
- implement a self-organising diagnostic capability with the ability to deploy secondary sensing.
- complete second stage of self-organising diagnostic capability: to distinguish between system failures (e.g. failure of electronics, comms.) and two different impact types (fast particle, slow, heavy impact)."

While there has been some departure from the specific requirements of the SoW, the general requirements have been satisfied. Developments to the CD system that will be reported here are as follows.

1. Development of a mobile robotic agent and its integration into the CD sensing and communication network.
2. Development of ultrasonic communication channels between the robot and the embedded agents of the CD that are in the immediate vicinity of the robot. This preserves the principle of agents communicating directly only with their immediate neighbours.
3. Implementation and demonstration of a fully self-organising inspection and diagnostic system with the ability to deploy secondary sensing using the robotic agent to perform this sensing function.
4. Demonstration of a self-organising diagnostic capability that can distinguish between system failures and two different impact types.
5. Incorporation of a video damage inspection capability on the robot, and a wireless communication facility. These are not yet fully functional: the hardware has been constructed, and the software designed but not yet fully implemented.
6. An active ultrasonic capability for secondary sensing has not yet been deployed, although the hardware infrastructure to do so is largely in place: fixed ultrasonic transmitters have been installed on the embedded agents, and the robot has an ultrasonic transmitter. This has not yet been done because, for practical reasons, the skin of the structure is not yet being subjected to real damage.

Structure of report

The next section provides an overview of the development and operation of the CD system as a whole, and this is followed by more detailed reports of the two major components of the system, the fixed structure with its embedded agents, and the mobile robotic agent. The overview is provided first in order to set the context for the subsequent component descriptions.

These sections are followed by a short outline of modifications to the system visualiser, a summary of what has been achieved and an indication of the next steps in this development program. An appendix lists the publications and reports that represent the written output of this program since its inception in late 2001.

2. System Overview

2.1 Introduction

The purpose of this section is to provide a high-level overview of the self-organising SHM test-bed and concept demonstrator (CD) that is under development at CSIRO, and which is the subject of this report (see [2, 3] for some background to this work). Further details of the various hardware and software components will be given in subsequent sections.

The system architecture is that of the complex multi-agent system [4], in which each of the semi-autonomous sensing agents contains a suite of sensors to enable it to gather data related to the state of the structure (generally, but not necessarily, referring to the agent's local region), a facility to perform some processing of this data, and the ability to communicate with neighbouring agents. These agents may also have the capability to perform active tasks (e.g. repair functions), or there may be other agents to perform these functions. Agents may be embedded in the structure, eventually being integrated into the materials, or may be mobile and free to roam regions of the structure.

The present CD system contains both embedded and mobile agents, and much of the recent work has been concentrated on integrating a single mobile robotic agent into the existing system of 192 embedded agents. The single robotic agent is the forerunner of a group of smaller robots capable of acting cooperatively to perform either sensing or repair tasks.

2.2 CD structure: embedded agents

The physical structure of the CD is a hexagonal prism with an aluminium skin (Figure 2.1). The skin consists of 48 aluminium panels (eight on each side of the hexagon), each of which contains four "cells". Cells are the fundamental building blocks of the system: they are the electronic modules containing the sensing, processing and

Figure 2.1: The hexagonal prism physical implementation of the test-bed, lying on its side with the end open to reveal the cellular internal structure of the electronics.

communication electronics. Each cell is an agent of the distributed multi-agent system. It communicates with its four immediate neighbours. The terms "agent" and "cell" are used interchangeably throughout this report. Each cell occupies an area of ~100 mm x 100 mm of the skin, mounted on the inside of which are four piezoelectric polymer (PVDF, polyvinylidene fluoride) sensors to detect the acoustic waves that propagate through the skin as a result of an impact. A fifth piezoelectric element, in this case a PZT (lead zirconate titanate) ceramic, is located in the centre of the square formed by the four PVDF sensors. This may be used as a generator of elastic waves, for two purposes: to actively probe the skin for damage evaluation; and to transmit signals for direct ultrasonic communication with mobile agents.

2.3 Damage simulation and initial sensor-based diagnosis

The present function of the CD system is to detect impacts and diagnose the level of damage to its operation. The software system has been designed to distinguish between four possible types of damage: hard impacts resulting in damage to the panel (a critical impact, requiring rapid attention by the robot, or mobile agent), lesser impacts resulting in damage that does not need immediate repair (a non-critical impact), failure of the communications channels between cells, and failure of the electronics within a cell. In order to avoid real damage, and consequently costly repair to the panels, the impacts used did not cause damage, and impact-generated damage was simulated – after a high-energy impact a red marker was stuck to the panel at the site of the impact. For lesser impacts a green marker was used.

An initial diagnosis of impact severity was made from the sensor outputs of the impacted cell (agent): a general discussion of the approach to damage diagnosis based on self-organisation is given in [5]. In this case self-organising maps (Kohonen neural networks [6, 7]) have been implemented to classify impact severity, distinguishing critical impacts from non-critical impacts. A previous discussion of the application of SOMs to the analysis of impact data is given in [8]. Electronic and/or communication failures, which are detected when a cell loses some or all of its communication capability, are distinguished from critical impacts that have damaged the electronics by the absence of an impact recorded by a neighbouring cell.

2.4 Robotic mobile agent

The deployment of mobile agents on or within a structure has the potential to provide an efficient means of carrying out active functions to enhance the capability of the embedded SHM system. While the eventual aim is for a swarm of small robotic agents, at this stage only a single, relatively large robot has been implemented. A critical aspect of the operation of mobile agents is that they are components of the multi-agent system: they have no central controller; they communicate with neighbouring agents (fixed or mobile); and their behaviour is determined by self-organisation.

The robot (Figure 2.2) has two legs and can move about the external surface of the CD skin rather like an inch-worm. It attaches itself to the skin using suction cups in both feet and moves either by sequentially stretching out and contracting in a linear motion, or by detaching one foot and pivoting around the other. Initially the robot will carry two small video cameras, one on each foot, which will send images back to

the network for damage evaluation. Other sensors may be included, such as an active ultrasonic transducer that can interact with the piezoelectric receivers embedded in the skin for ultrasonic damage evaluation.

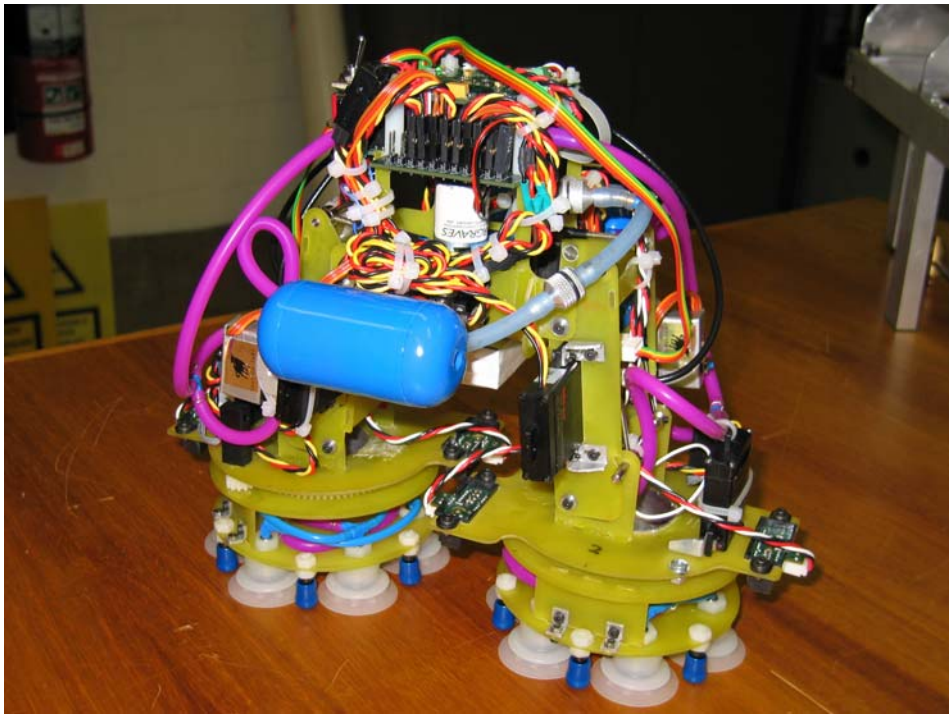

Figure 2.2: The robot on a horizontal wooden bench.

The size of the robot is such that each foot occupies a significant fraction of the area monitored by a single embedded agent. Each step moves a foot to the region of a different agent, and the robot attempts to position the moving foot at the centre of the target agent. The robot navigates entirely by information it receives from embedded agents in the structure. The methods by which this motion and navigation are achieved will be outlined below, and discussed in more detail in subsequent sections.

2.5 Communications

The primary mode of communication between the robot and the underlying embedded agents is via the acoustic coupling provided by the aluminium skin of the CD structure. An ultrasonic transducer, described in more detail in Section 4, is mounted in each of the robot's feet: when a foot is attached to the skin, two-way communication can take place between its transducer and the sensors of the agent in whose vicinity the foot is placed. At this stage, only three forms of communication take place: firstly there is dialogue that enables the robot to re-position its foot close to the centre of the square bounded by the four PVDF sensors of the embedded agent. This ensures that the robot's next step will begin from a known position, and prevents cumulative stepping errors.

The second form of communication provides navigation information from the embedded agents, which enables the robot to decide in which direction to take its next step. The nature of this information will be described below, but essentially it allows

the robot to decide to move to sites with different levels of damage, or to a docking station for a battery re-charge (although this facility has not been implemented yet).

The third form of communication allows the robot to send its sensor data back to the CD system. Implementation of this is not yet complete, but it will utilise a wireless channel that can operate in parallel with the ultrasonic communications. Initially, the wireless receiver will be located in the visualiser computer, an external agent that displays the state of the system on a monitor, but later there will be a number of wireless transceivers on embedded agents around the structure.

2.6 Robot navigation

The requirement is for each embedded agent to contain information that will enable the robot, once it has placed its foot on the agent, to determine its next move. The method by which this is achieved in this case is based on a distributed dynamic programming algorithm employing a gradient field set by each impact cell [9]. A gradient is initiated by an impacted cell, and propagated successively to neighbouring cells. Figure 2.3 shows the gradient produced from two non-critical impacts, located in the black cells. The gradient value at each cell is indicated by the shade of green. White cells are those with which the robot cannot communicate. This may be due to electronic failure or, as is the case here, sensors have not yet been fitted.

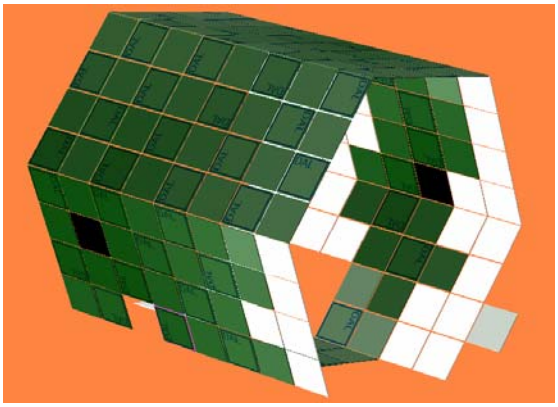

Figure 2.3: An image of the gradient produced by two non-critical impacts that occurred at the black cells. The squares indicate the cells on the surface of the hexagonal prism test-bed. The gradient values are shown as shades of green, with a higher gradient being darker. The white cells are those with which the robot cannot communicate. Absent cells in the image are those that have been physically removed.

In addition, the gradient-based algorithm incorporates multiple fields to enable an optimal prioritisation to account for the relative importance of visiting severe impact sites rapidly versus the time saved by visiting nearby sites. A gradient field is also established to guide the robot back to its base, in the event that battery recharging is required. Dynamic modifications of the gradients occur when new impacts are recorded, or when inspections and/or repairs of impact sites are completed.

When the robot foot is placed on a cell, the robot receives the value of the gradient (or gradients for multiple fields) at that cell, and for the neighbouring cells, and from these it determines the direction of its next move. This decision is based on its

priorities (e.g. always go to docking station if battery charge is less than some threshold; otherwise, always go to the nearest high severity impact; if no high severity impact, go to nearest impact or follow steepest gradient).

This navigation decision by the robot is based purely on information obtained from the local agent. The gradients are a self-organised response to the detection of an impact by a local agent. *Therefore, the robot movement is a self-organised response of the system of agents to the impact.*

Initially, the function of the robot is to obtain additional sensor data that can be used for damage diagnosis, i.e. it is an active inspection/monitoring agent. The robot is fitted with a video camera on each foot to enable direct inspection of damage, but it could equally well be fitted with other sensors (e.g. for active ultrasonic inspection or eddy current sensing) as well as or instead of the video camera. Eventually the robot (or robots) could also perform repair functions. The ultimate vision for a system such as this is to have swarms of very small robotic agents that can work cooperatively to perform both inspection and repair tasks.

3. CD Embedded System Components: Hardware and Software

3.1 CD architecture and hardware

The physical structure of the test-bed is a hexagonal prism with an aluminium skin (Figure 3.1), with the prism ~ 1 m in height and ~ 1 m across the hexagonal cross-section. Mechanical, electronic and software aspects of the CD/test-bed are contained in our most recent report [1].

Figure 3.1: The hexagonal prism physical implementation of the test-bed, lying on its side with the end open to reveal the cellular internal structure of the electronics.

The initial goal of the test-bed is to detect and characterise impacts on the skin, and to form a diagnosis of the accumulated damage. The skin consists of 48 aluminium panels (eight on each side of the hexagon), each of which contains four "cells". Cells are the fundamental building blocks of the system: they are the electronic modules containing the sensing, processing and communication electronics. Each cell is an agent of the distributed multi-agent system. It communicates with its four immediate neighbours.

Each cell occupies an area of ~100 mm x 100 mm of the skin, mounted on the inside of which are four piezoelectric polymer (PVDF) sensors to detect the acoustic waves that propagate through the skin as a result of an impact. Thus the complete test-bed contains 192 cells. One of the panels, and its four cells, is shown in Figure 3.2.

The cell electronics are constructed as two sub-modules, each 80 mm × 80 mm and mounted directly on top of each other as shown in Figure 3.2. One of the sub-modules, called the network application sub-module (NAS), contains the communications and processing hardware, while the data acquisition sub-module (DAS) contains the analogue electronics and digitisation hardware specific to the attached sensors. A benefit of this division is that the NAS is flexible enough for almost any SHM sensor network application, and only the DAS needs to be changed to accommodate the different sensors that may be required in different applications. Further details of the electronics can be found in [1, 10].

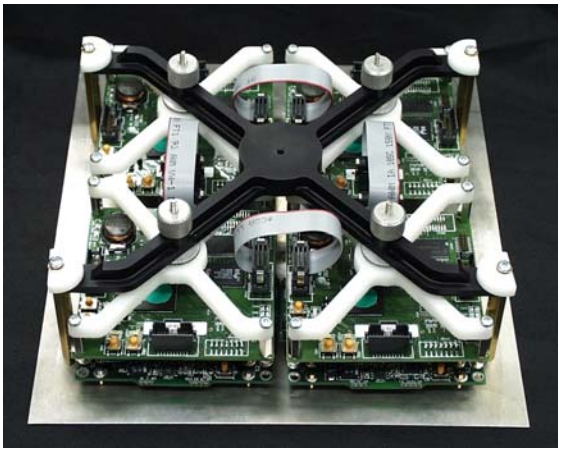

Figure 3.2: Aluminium panel containing four cells. Each cell consists of a data acquisition sub-module (DAS) below a network application sub-module (NAS). Each cell is connected to its four immediate neighbours, via the ribbon cables that can be seen in the photograph, to form a square network array.

3.2 Piezoelectric sensors

The layout of the array of sensors on panels of the Concept Demonstrator has been described in detail in previous reports [1]. The aluminium panels have dimensions 200 mm x 220 mm x 1 mm. Each panel has bonded to one side an array of transducers. This consists of sixteen PVDF discs (110 μm in thickness, 2.5 mm in diameter, coated on one side with silver ink, while the other side is bonded to the aluminium) and four PZT discs (0.5 mm in thickness, 2.5 mm in diameter, with fired-on silver electrodes on the two flat faces). These are bonded to the panels with Araldite 2015 epoxy in the arrangement shown in Figure 3.3.

In the fully populated Demonstrator there are 48 panels, and therefore 768 PVDF and 192 PZT transducers. A photograph of a panel is shown in Figure 3.4.

The size and positioning of the sensors has been discussed at length in previous reports. It is sufficient to say here that the sensor size was largely determined by three factors: the requirement early in the project of being able to determine the location of impacts by triangulation of pulse-echo signals with a precision of the order of a few millimetres, the necessity of having the PVDF sensors large enough to generate a measurable voltage, and yet small enough to avoid the effects of phase cancellation if the size of the sensor becomes large compared with the wavelength of the acoustic waves.

Bonding transducers consistently to an aluminium panel requires care – the layer of glue needs to be as thin as possible without compromising the strength of the bond, little or no excess adhesive around the circumference of the transducer is desirable, and the sensors must not be accidentally shorted electrically to the aluminium panel, as can sometimes happen with the thin PVDF sensors that are punched out of a large, coated sheet of PVDF. While this is largely a matter of practised skill, every sensor was checked prior to the attachment of the electronics to the panel.

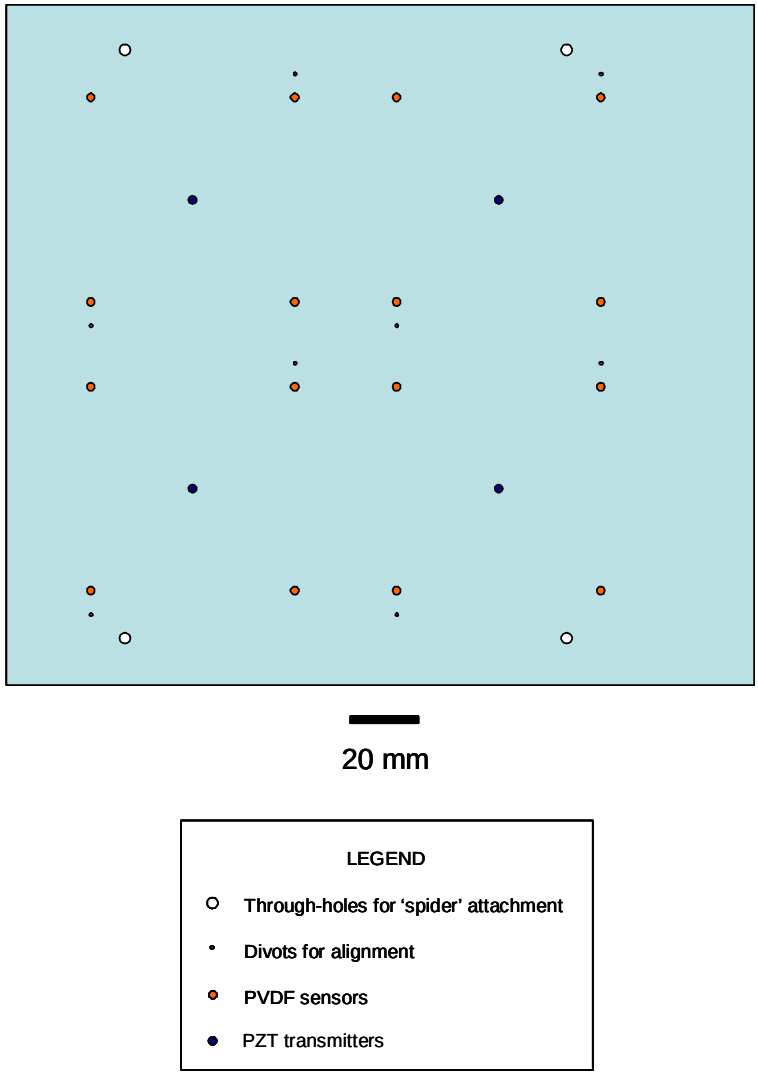

Figure 3.3: Arrangement of sensors on aluminium panel.

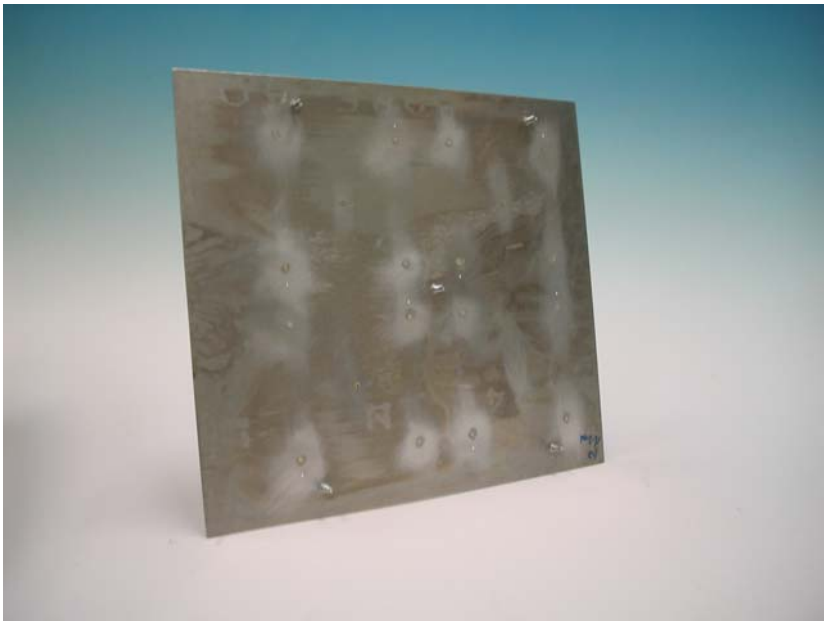

Figure 3.4: Photograph of aluminium panel with sensors attached.

A simple and practical method of checking PVDF sensor function and output was to strike the panel and monitor the sensors' voltage output. For this purpose a jig was constructed (Figure 3.5), consisting of a U-shaped holder for the panel, a pendulum for striking the panel, and a set of electronics (on a board with spring-loaded pins to connect to the sensors as does the standard DAL board of the CD) to measure the voltages from each group of four PVDF sensors. The U-shaped holder was built in such a way as to allow the panel to be accurately positioned to have the pendulum strike at the centre of each cell (i.e. at the centre of each group of four PVDF sensors). While some variation in the output of these thin sensors is expected, this method easily identified poorly bonded sensors which could be removed and replaced with a new sensor. The PZT sensors were checked individually on a gain-phase analyser.

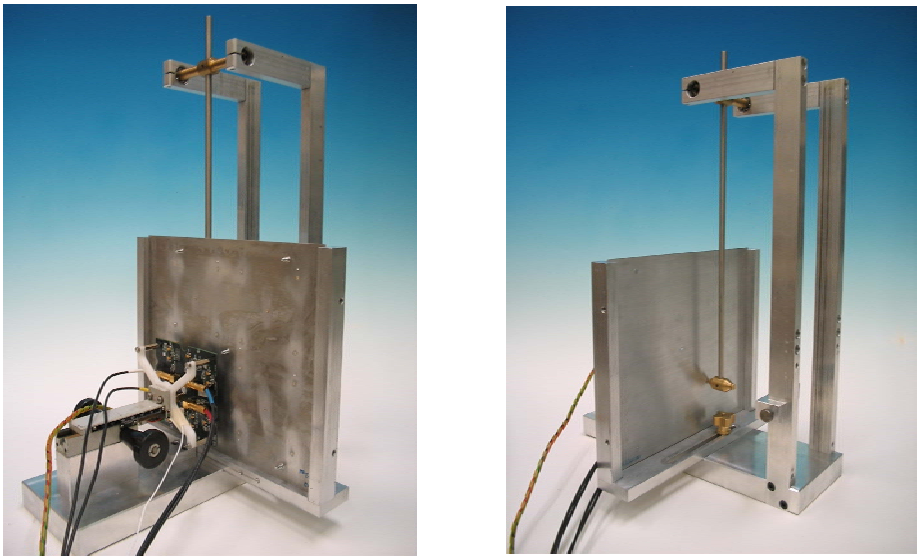

Figure 3.5: Photographs of pendulum device for testing sensor function.

Originally the role of the PVDF transducers was simply to detect, characterise and locate impacts on the aluminium skin. Each group of four PVDF sensors in each cell would detect elastic waves resulting from an impact travelling through the aluminium plate, and use this information to determine the location and severity of the impact. This role has now been expanded: these transducers also act as the receivers of communication signals from the mobile robotic agent.

The PZT transducers were to fulfill a number of functions. They were to be used as transmitters to send ultrasonic waves through the panel for subsequent detection by the PVDF sensors before and after impacts on the panels. This allows firstly the calibration of the PVDF sensors (or at least a measurement of their sensitivity), and secondly the possible determination of the state of damage of the panel after the impact has occurred. They were also envisaged to be transducers for bi-directional ultrasonic communication through the panel, but because of a deficiency in the original design of the DAL boards (for which such communications were not anticipated) this central PZT is used only as the communications transmitter.

3.3 Modifications to the CD panels

The present incarnation of the mobile robotic agent attaches itself to the surface of the CD by two sets of suction feet. In order to have each foot reliably and firmly attached to the aluminium panels it was decided to modify the attachment fixtures that held the electronics in place. The round socket-head cap screws were replaced by flat self-clinching studs. This made the outside surface of the panels smooth and flat. While this proved adequate to maintain a vacuum sufficient to allow attachment of the robot, a further modification for better orientation of its feet was found to be required, which additionally ensured excellent attachment.

The robot uses two infrared rangefinders on each of its feet to correctly orient the feet for placement on the panel. While the servo motors align the feet approximately, a feedback loop using the rangefinders constantly makes fine adjustments to the foot's orientation as it approaches the surface of a panel, and does so until the suction has taken full effect, at which time the rangefinders are turned off. The rangefinders also prove useful for preventing movement of the robot that may cause damage to its feet. The standard rolled finish of the aluminium panels proved to be a poor reflector for the rangefinders, and so several surface treatments for the panels were considered and tested. The easiest and most effective was simply to apply a matt white adhesive-backed plastic film, similar to the type of plastic used to cover books. With care, this could be applied without bubbles or any other surface imperfection. It was important, however, to allow the adhesive to cure fully before applying any de-bonding force, such as the robot's suction feet. The adhesive does not reach full strength until a few hours after application. This plastic film did dampen the acoustic waves slightly, but there was still plenty of signal for triangulation and communication purposes. Another benefit of the film was to dampen reverberations in the panel. As communications between the robot and the CD relies on the transmission and reception of a continuous train of phase-shift-keyed tones, the reduction of reverberations in the panels allowed much cleaner signal transitions, and therefore error-free information transmission.

3.4 Impacts and simulated damage

As outlined in Section 2, the CD's present function is to detect impacts and diagnose the level of damage to its operation. The software system has been designed to distinguish between hard impacts resulting in damage to the panel (a critical impact, requiring attention by the robot, or 'mobile agent'), lesser impacts resulting in damage that does not need immediate repair (a non-critical impact), and electronics and/or communications failures not caused by an impact. Real damage was avoided, and impact-generated damage was simulated by different coloured markers stuck to the panel at the impact site: a hard impact was indicated by a red marker and a lesser impact by a green marker. This allowed the secondary inspection system (using the robot with a small video camera and simple frame-grabbing software to determine the colour of the marker and hence the 'severity' of the damage, described later) to locate and diagnose the damage for comparison with estimates made from the piezoelectric sensor data.

The 'hard' and 'soft' impacts could be applied by simply rapping the outside surface with a hard object, and just guessing the energy with which this was done. Alternatively, and much more usefully in the early stages of the project when some degree of calibrated impact was necessary for debugging the algorithms, hits of fixed (high or low) energy needed to be available. To do this the pendulum apparatus used to test the performance of all the PVDF sensors immediately after bonding to the panels and before the installation of electronics, was modified. Shown in Figure 3.6 this modified, hand-held version could be set to deliver any energy impact up to the highest potential (and hence kinetic) energy obtainable from the pendulum, repeatably and reliably. It was held against a panel (using its three felt-covered stand-offs), and the pendulum drawn back to the desired and repeatable height for striking the panel with the selected energy.

Following such an impact, the PVDF sensor data was used to estimate the position (using the triangulation method outlined in Report 5 [1]) and severity (using self-organised maps, as described in next sub-section) of the measured impact.

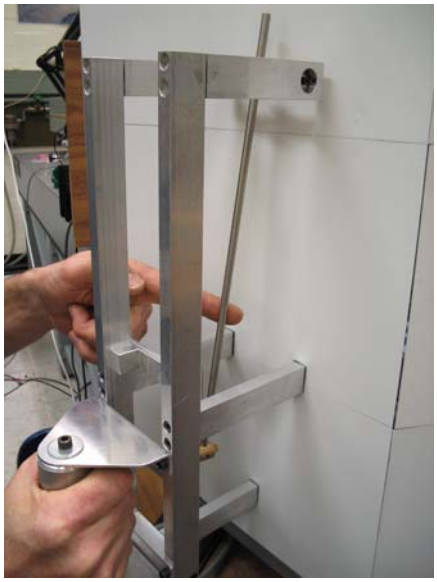

Figure 3.6: Hand-held apparatus for delivering fixed energy impacts.

3.5 Impact signals and sensor-based diagnosis: use of self-organised maps (SOMs)

A general discussion of the approach to damage diagnosis by self-organisation is given in [5]. In this case self-organising maps (Kohonen maps [6, 7]) have been implemented to classify impact severity, distinguishing critical impacts for which the skin has been ruptured, from non-critical impacts. A previous discussion of the application of SOMs to the analysis of impact data is given in [8]. Electronic failures, which are detected when a cell loses its communication capability, are distinguished from critical impacts that have damaged the electronics by the absence of an impact recorded by a neighbouring cell.

The aim of this signal-based diagnosis is to identify high- and low-severity impacts in different regions of a panel: specifically, whether an impact has occurred within the cell that has recorded the signals or within one of the other three cells of the panel. In general, a cell's sensors will detect an impact that occurs anywhere on the panel on which the cell is located, but usually not if the impact occurs on another panel. The diagnosis should be able to unequivocally identify on which cell the impact occurred (even if that cell has been damaged to the extent that it can no longer communicate), an approximate position within that cell, and whether the impact was of high or low severity.

A self-organised map (SOM) [4, 5] is an unsupervised learning algorithm that projects high-dimensional sets of data, in this case time series of sensor data, onto a two-dimensional grid. The form of the projection is determined automatically from the data. Each node of the grid corresponds to a model M_i that is represented by a vector (the model vector or codebook vector) that has the same dimension as each data vector. These vectors, and the models they represent, develop a stable pattern as training of the system progresses. The SOM algorithm ensures that similar vectors are located close to each other on the grid, leading to spatial self-organisation of the information in the grid.

Briefly, the SOM algorithm works as follows. The codebook vectors are initialized with random values. If $\boldsymbol{x} = (x_1, x_2, \ldots, x_n)$ is a data vector (a vector formed from the time series of the digitised sensor output), the first task is to find the codebook vector to which $\boldsymbol{x}$ is most similar. In cases like this, the closest vector (the 'winner') is defined as that with the minimum Euclidean distance from $\boldsymbol{x}$. If the codebook vector on the i^{th} node of the grid is $\boldsymbol{m}_i = (m_{i1}, m_{i2}, \ldots, m_{in})$, this is defined as:

$$c = \arg\min \{ \| \boldsymbol{x} - \boldsymbol{m}_i \| \}$$

where the minimum is over i values between 1 and N, the number of nodes of the grid.

When the winner ($\boldsymbol{m}_c$) has been identified, codebook vectors in its immediate neighbourhood, including the winner itself, are up-dated to make them closer to $\boldsymbol{x}$. This is the origin of the spatial organisation of the map. It is done by up-dating the codebook vectors as follows:

$$\boldsymbol{m}_i(k+1) = \boldsymbol{m}_i(k) + \eta(t) \, . \, h_{ci}(t) \, . \, (\boldsymbol{x}(k) - \boldsymbol{m}_i(k))$$

where $k = 1, \ldots, K$ is the data vector index, $h_{ci}(t)$ is a neighbourhood function that defines the spatial region about $\boldsymbol{m}_c$ that will be modified – it may be a Gaussian or a square function, for example – and $\eta(t)$ is a learning rate function that is used to reduce the rate of modification of the vectors with increasing number of learning iterations t. At each iteration t, there are K updates: one for each training vector $\boldsymbol{x}(k)$. The indices k and t indicate that the set of K training vectors is used for $t = 1, \ldots, T$ training iterations (epochs). Commonly, the learning rate is reduced and the neighbourhood function is narrowed as the number of iterations increases.

Training of the self-organised maps (SOMs) was done on a single panel, using hard and soft impacts produced by the apparatus shown in Figure 3.6 at a number of locations within each of the four cells on the panel.

The initial aim was to use the SOM to identify the following four conditions:

- A soft impact occurred within the cell.
- A hard impact occurred within the cell.
- A soft impact occurred outside the cell.
- A hard impact occurred outside the cell.

If this diagnosis can be made for each cell on a panel that has suffered an impact, then the impacted cell can be identified unambiguously.

For a particular cell on the panel, 100 soft and 100 hard impacts at random positions within the area of the cell were sampled. These samples covered most of the cell's area thoroughly, giving roughly 3 impacts per square cm within the cell. Further, for impacts outside the cell, 100 soft and 100 hard impacts were sampled from the areas of the 3 remaining cells.

An impact event is recognised when a sensor signal threshold is exceeded. The cell DAL board then acquires 256 samples from each of its four sensors. It cannot be assumed that any of the signals reflect an accurate time of arrival of the impact pulse, because of the use of a constant threshold. In order to reduce the memory requirement for each SOM, and to maximise the size of the SOM array that can be stored, the string length was reduced to 64 by 4-point sub-sampling of each 256-sample signal. A data input vector $\boldsymbol{x}(k)$ used for training the SOM consists of the 64-point data string from each of its four sensors in the relevant cell. This allowed a 10 x 10 SOM array to be stored in binary format in 50 kB of flash memory on each agent.

For the purpose of forming the data vectors the four sensors were ordered based on time of arrival of the signal, because the training set is taken for one particular orientation and the cells on the CD are not always in the same orientation. This makes the SOM orientation-independent, but at the expense of geometric information about the direction of position at which the impact occurred. This is an issue that can be re-examined later.

In order to improve the efficiency and effectiveness of the learning process, the 10 x 10 SOM was trained in an unconventional way. The 10 x 10 array was divided into four 5 x 5 arrays, with each of these smaller SOMs assigned to learning only one of the four conditions, listed above, that are to be identified. Each of the 5 x 5 blocks was then trained separately using the sub-set of the training signals that corresponded

to the particular block, e.g. the 5 x 5 block assigned to soft impacts within the sensing cell was trained using only the signals produced by soft impacts within the sensing cell. The boundaries of the 5 x 5 blocks were wrapped around (top to bottom, left to right) to provide continuity at the edges. The four 5 x 5 blocks were then assembled into a single 10 x 10 SOM array, as illustrated in Figure 3.7. While the conventional approach to SOM formation is purely unsupervised data-driven learning, the present approach may be considered to be adding an element of supervision to the learning process.

Training of each of the four 5 x 5 SOMs was controlled by the following parameters:

- Neighbourhood function $h_{ci}(t) = 1$, i.e. a constant over the 5 x 5 block, independent of location in the array and training iteration number. This was a simplifying assumption resulting from the small size of the blocks.
- The maximum number of training iterations (epochs) was $T = 2500$.
- The learning rate function $\eta(t)$, which is required to decrease with increasing training iteration number, was taken as:

$$\eta(t) = 0.10 \cdot e^{(-t/T)}$$

The trained vectors thus produced were stored as a SOM on each of the agents on the CD. When identifying the type of impact, the four smaller SOMs are evaluated as a single 10 x 10 SOM.

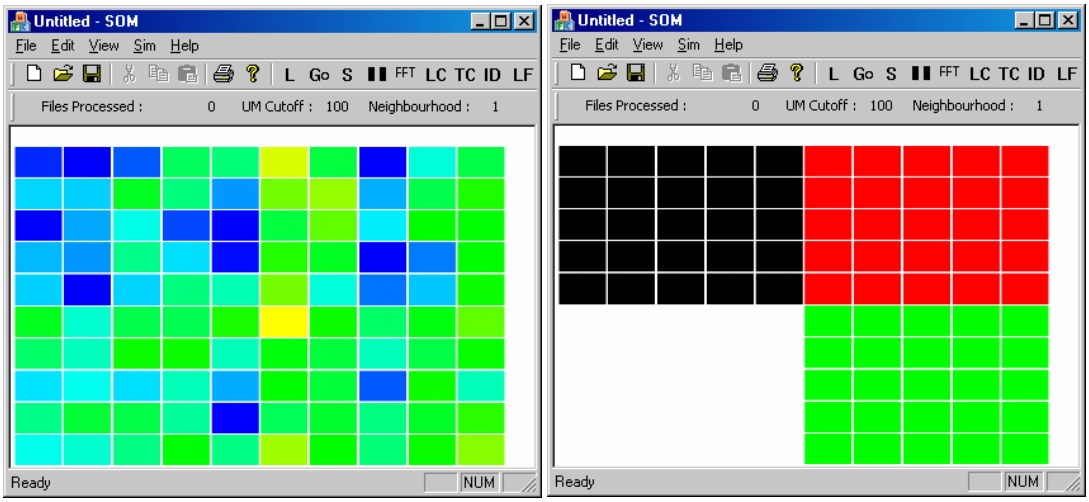

Figure 3.7: (Left) U-matrix visualisation of the four 5x5 SOMs as one 10x10 SOM. The U-matrix is a visualisation method that assigns a colour to a cell according to the (Euclidean) difference between adjacent codebook vectors. It is useful for identifying emergent clusters of similar vectors, but it is unlikely that identifiable clusters will appear in small SOMs such as these.
(Right) Structure of the 10 x 10 SOM in terms of the 5 x 5 blocks. The black area represents the learnt vectors for hard inside impacts, white for soft inside impacts, red for hard outside impacts and green for soft outside impacts.

In order to evaluate the accuracy with which signals can be identified with the resulting SOM, only half of the training signals were used to train the SOM and the other half were used evaluate the accuracy of recall and precision. Several separate

runs showed a consistent accuracy of ~ 93 ± 1%, independent of the signal trigger threshold over a range from ~ 1.2% to ~ 3.6% of the maximum signal amplitude. Given that each impact is detected by four cells, the probability of an incorrect assignment of an impact location is very small.

For greater accuracy, the SOM on each agent could have been trained separately – this would have individualised the SOMs to take account of the inevitably different sensitivities of the sensors in different cells. Ultimately, it is expected that SOMs will learn on-line, so even if they are initially trained with a common data set, the subsequent learning will develop a set of individualised SOMs.

So far the SOM-based impact diagnoses have proved to be highly reliable, but tests to date have all been carried out using the same impact mechanism (i.e. the apparatus of Figure 3.6) with which they were trained. It remains to be seen whether the SOMs will retain this high accuracy with impacts of different origin but similar spectral characteristics.

3.6 Self-organised robot guidance: gradient field algorithms

There are two related methods by which navigation of the robot may be achieved, directed by self-organisation. Firstly, algorithms based on ant colony optimisation (ACO) [9, 11, 12] have been developed in earlier work to link sub-critical impact locations by a simulated pheromone trail and a dead reckoning scheme (DRS) that form a minimum spanning tree [9]. The decentralised ACO-DRS algorithm has low communication cost, is robust to information loss within any individual cells, and allows navigation around critically damaged regions in which communication capability has been lost. An example of a self-organised network connecting a number of sub-critical impact sites, which was produced by this algorithm, is shown in Figure 3.8.

An alternative scheme evaluated in [9] is a distributed dynamic programming algorithm, employing a gradient field (GF) set by each impact cell.

While the concept of the robot following the self-organised pheromone trails produced by ACO is appealing, there is a trade-off between the low communication cost of the ACO-DRS algorithm and a better quality of the minimum spanning tree approximation computed by the gradient-based algorithm [9]. The GF algorithm also ensures that each cell in the system has a valid gradient value: the ACO-DRS algorithm does not guarantee a pheromone value in every cell.

The approach that has been implemented is the gradient field (GF) algorithm. The basic principle of gradient propagation is very simple. All cells are initiated with a high value for the gradient. When a cell detects an impact, its gradient value is set to zero. At each time-step, each cell compares its gradient value with that of each of its neighbours. If a cell finds a neighbour with a gradient value lower than its own, ittakes this value plus a pre-defined increment (say δg) as its new gradient value. This is repeated at subsequent time-steps until a stable gradient field is produced over the whole array of cells, i.e. the gradient produced by the impact propagates throughout the system of agents. It is independent of the number of neighbours each cell has, so

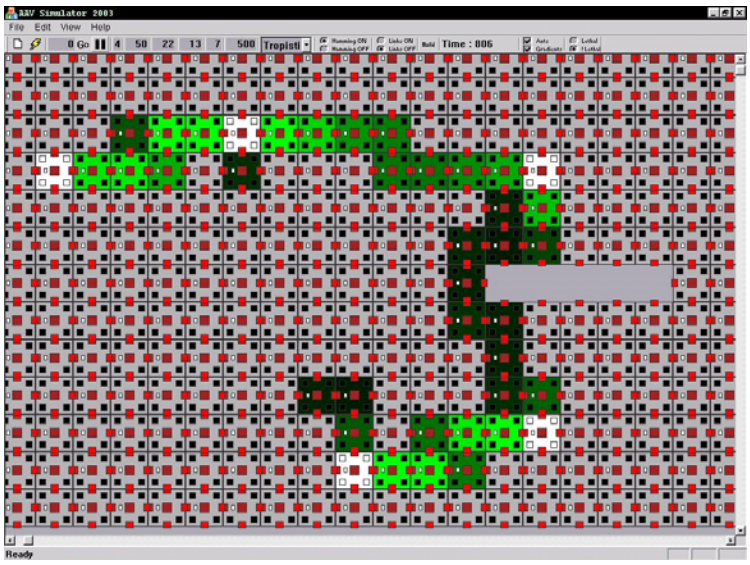

Figure 3.8: Computer simulation of a reconfigurable, self-organised impact network produced by the decentralised ACO-DRS algorithm. White cells are those that have received a sub-critical impact. Green cells are those that have a pheromone level greater than a threshold value, with lighter shades indicating a higher concentration. See Report 4 [13] for further details.

it is robust in the presence of failed cells. Multiple impacts produce a gradient field with multiple minima (see, for example, Figure 2.3), analogous to a topographic map of a surface with minima that correspond to impact sites.

If impact damage is repaired, or inspection shows that it can be ignored, the gradient field can be readily modified to reflect this change. This is achieved as follows. The previously impacted cell resets its impact flag, and then increases its gradient to a value that is a small increment greater than that of its neighbour with the lowest gradient value. All cells then check all their neighbours. If a cell has a gradient value lower than those of all its neighbours, and it hasn't been impacted, its gradient value is incremented by δg. Repetition of this algorithm removes the minimum associated with the repaired cell. This action is analogous to one of several separated weights being removed from the surface of a trampoline.

To allow the robot to identify and respond to different classes of event, a gradient field could be set up with minima of varying depths corresponding to the severity/importance of each event. However, steps would have to be taken to avoid the complication of the robot being attracted to nearby minima in preference to more distant but perhaps more important minima. A simple solution to this problem is to model different classes of event on separate gradient fields, and set all minima on a single gradient field to be equally important (deep). The multiple classes of gradient field may all model different routes to their respective sites of importance. The robot may then respond to events in order of priority by exhausting one class of gradients before switching to a gradient class with lower priority, and so on.

This multiple gradient field solution has been implemented and currently the robot responds to four classes of gradient/event:

- High severity impacts, representing critical damage – perhaps penetrating impacts and/or cell destruction.
- Low severity impacts, representing non-critical damage – non-penetrating impacts and damage which does not affect system behaviour.
- Damage not caused by an impact, such as communications and/or electronic failure.
- Dock. Models the shortest route to the robot's docking station for re-charging, downtime, etc.

The robot can adopt different criteria to prioritise the order in which it visits impact sites. A possible modification to the behaviour outlined above would be to allow the robot to visit sites of lesser importance if they are on or near the intended path to a site of greater importance.

The major benefits of this algorithm are its stability and its fast dynamic response. A mobile agent in the system may determine the direction of the shortest path to the nearest impact location through interrogation of the local cell group. This is analogous to a ball on a slope; the ball need not know where the bottom of the hill is to know which way to roll – simply knowing the gradient of its local piece of hill is sufficient. In the GF algorithm, the shortest distance to an impact location can be determined simply from the value of the gradient in a particular cell, since each cell increments the lowest gradient value of its group of neighbours by one unit, δg.

3.7 Communication with the robot: distinguishing impact and communications signals

Communication between an embedded agent (cell) and the robot utilises ultrasonic signals propagated through the aluminium skin of the cell. The robot transmits and receives ultrasonic signals using transducers mounted in the centre of each foot. Further details about the robot's transducers and the communications sequences are given in the next section. The agent transmits via the PZT element in the centre of the cell, and receives signals through the same four PVDF elements that it uses for impact detection.

A communication is initiated when the robot places one of its feet on the region of skin monitored by agent *A* (say). In order to initiate a communication sequence, the robot then transmits a tone burst from the ultrasonic transducer in this foot, which consists of 5 cycles at a driving frequency of 400 kHz. This corresponds approximately to the lowest order radial mode of the robot transducer disc, and it excites the A_0 guided mode of the aluminium skin. The agent *A* distinguishes this signal from that due to an impact on the basis of its spectral content.

A 5-cycle tone burst has a spectral width of ~ 20% of the centre frequency, and this will be increased by the spectral response of the transmitting and receiving transducers. Nevertheless, it is expected to have significantly different spectral characteristics than an impact-generated signal. At this stage the agents use a simple combination of the responses of two band-pass filters with different cut-off frequencies to distinguish impact from communications events, but more sophisticated processing could be readily implemented.

A communications sequence is completed when the agent *A* receives a specific acknowledgement signal from the robot. The issue of an impact that occurs during a communications sequence has not yet been dealt with: at this stage it may result in a corrupted communications packet, but will not otherwise be recorded. This is not an urgent issue for the present system, since the robot foot would shield the cell during a communications sequence, though an impact might damage the robot. However, it raises potential issues of impact damage to the robot, and impact damage to the cell when much smaller robots are in use.

More details of the cell-robot communications are given in the next section.

4. The Mobile Robotic Agent

4.1 Development of the mobile agent

An important feature of the CD system is an ability to support mobile (robotic) agents that can roam the exterior surface of the test-bed, communicating with the fixed agents embedded in the underlying structure. The function and operation of such an agent will be described in this section, and it should be emphasised that it is not controlled centrally, but cooperatively with the network of fixed local agents with which it communicates.

It should also be emphasised that the system described here is no more than a test-bed, whose primary purpose is for investigation of the practicality of the self-organised complex system approach to damage diagnosis and response. Thus, details of the specific hardware implementation (such as the use of suction for the robot's attachment, which is obviously inconsistent with a space-based application) are not considered to be important at this stage. While the present implementation of the robot is bulky and represents a single point of failure, the eventual aim is to develop a swarm of very small robots that can perform internal or external tasks cooperatively. The work described in this report represents a first step towards that ultimate goal.

Why is a robotic agent needed? When sensing impacts using passive sensors, the information received may be insufficient to characterise the damage, and where damage is detected it may need to be repaired. One approach to obtaining additional damage data, and to providing a crude repair capability, is the development of a mobile robot that can move around the outside skin (see Figure 4.1).

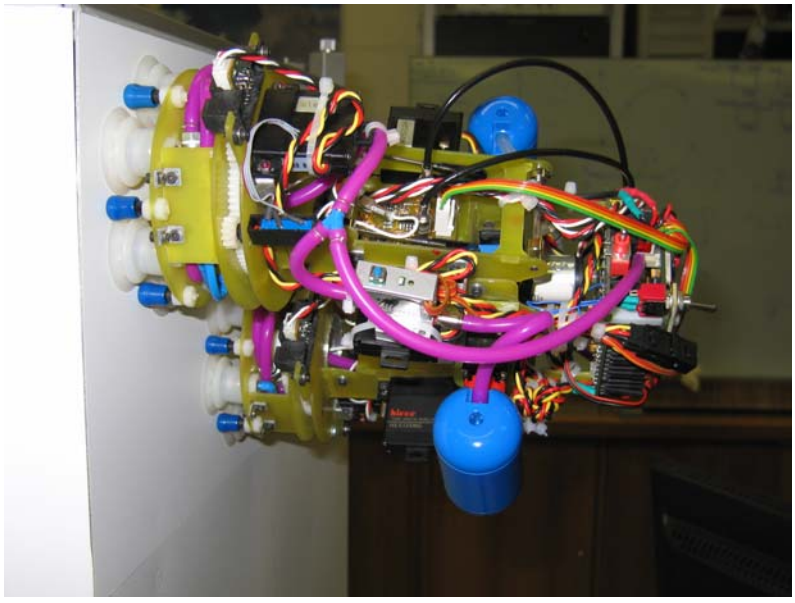

Figure 4.1: The robot on a vertical face of the test-bed.

The robot moves rather like an inch-worm, with its design based on an articulated box section with six degrees of freedom. The joints are driven by commercial model aircraft servos and have no position feedback to the controlling processor. The robot

is equipped with six suction cups on each of its two feet, and a pneumatic system with a variable speed vacuum pump and electrically controlled valves that allow it to selectively attach and detach its feet to and from the surface. To allow the robot to find the surface and attach to it reliably there are two optical rangefinders on each foot that measure the distance to the surface and the angle of the surface. A lithium polymer battery supplies power to the robot for approximately 30 minutes of operation before recharging is necessary.

The robot attaches itself to the skin using suction cups in both feet and has two modes of locomotion. The first mode is very much like an inch-worm: to move forward the robot alternately stretches out and contracts whilst detaching and attaching its feet in sequence. The second mode requires the robot to detach one foot, pivot 180° around the other (attached) foot and then reattach the first. It can change direction by pivoting through any angle up to 360°. Initially the robot will carry two small video cameras, one on each foot (Figure 4.2), which will send images back to the network for further analysis. In the future other sensors may be included, such as an active ultrasonic transducer that can interact with the piezoelectric receivers embedded in the skin for ultrasonic damage evaluation.

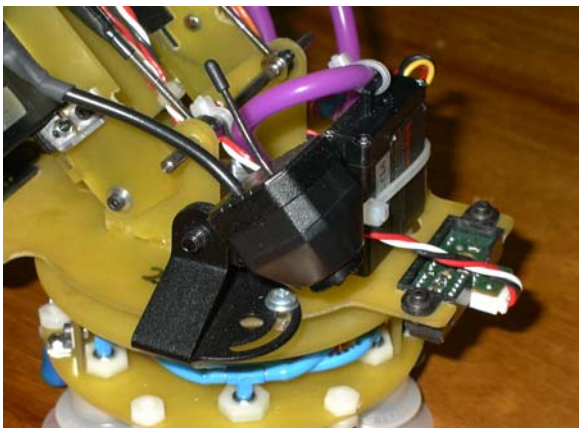

Figure 4.2: Close-up of one foot of the robot, showing the inspection video camera (centre, foreground), and one of the optical range finders (right).

Figure 4.3: The ultrasonic transducer mounted in the centre of the robot foot.

The robot communicates with the fixed agents in the network using piezoceramic (PZT) ultrasonic transducers in both feet (Figure 4.3) to pass messages through the aluminium skin to the underlying cells. The development and performance of these transducers will be described below. The fixed agents receive messages via the four piezoelectric polymer sensors that are used for detecting impacts. A fifth transducer, in this case a piezoceramic, has been added at the centre of each cell for transmission of messages from the cell to the robot.

4.2 Ultrasonic sensor development

On each foot of the robot is mounted an ultrasonic transducer used for communicating with the cell on which it is placed, and for indicating to the cell its foot position. This transducer had several stages of development in order to achieve a small, robust and reliable high-output device. The first version of this device is described in detail below. Subsequent improvements, based on a similar method of construction, are then outlined.

The transducer's active element originally consisted of a 2.5-mm diameter, 0.5-mm thick PZT ceramic disc, with fired-on silver electrodes on each flat face. One side of this disc is bonded with electrically conductive (silver-loaded) epoxy to a piece of copper cut to cover the whole of the face of the disc and provide an earth return path. To the other side of this copper sheet a wear plate of alumina (aluminium oxide) is bonded, protecting the transducer from abrasion from the aluminium skin of the CD. The back side of the PZT disc is the active electrode (Figure 4.4).

This assembly fits into a cylindrical Ertalyte (PETP, a tough, easily machineable thermoplastic polyester with very good dimensional stability) housing, resulting in a small, wear-resistant transducer with electrical leads ready to be connected to the nearby electronics (Figure 4.5).

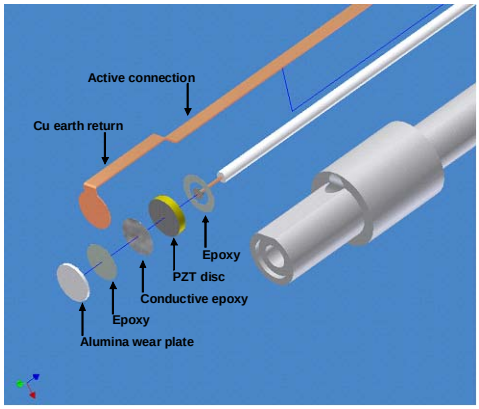

Figure 4.4: Exploded view of transducer

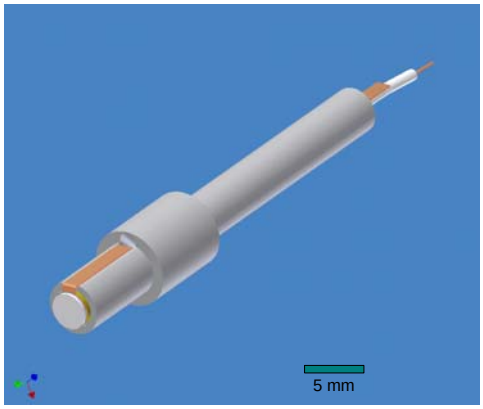

Figure 4.5: Assembled transducer

This transducer housing is spring-loaded and inserted into a threaded rod, with two thumb-screws on the outside enabling the position of the transducer with respect to the foot-plate of the robot to be varied (Figure 4.6). This construction allows us to

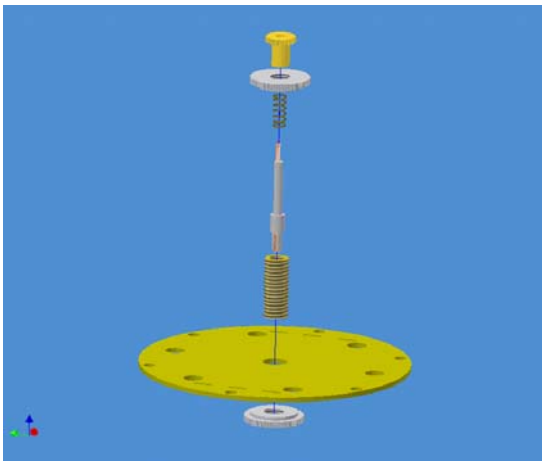

Figure 4.6: Exploded view of transducer mounted in a spring-loaded fitting that allows the position of the transducer to be adjusted relative to the CD surface.

vary the force with which the transducer is pressed against the skin, maximizing the ultrasonic coupling, and minimizing wear. Note that the design aim was to use dry coupling, i.e. to couple this transducer to the aluminium skin without any fluid couplant. One of these transducers is shown in the centre of one of the robot's feet (Figure 4.7).

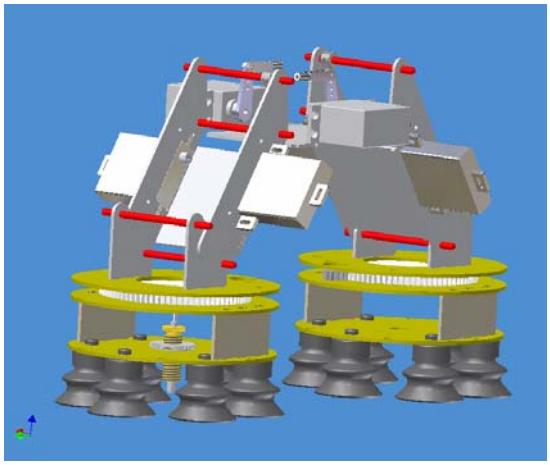

Figure 4.7: Schematic of the major structural components of the robot showing the ultrasonic transducer attached to one of its feet.

For a number of reasons, including the need to supply more energy through its dry coupling to the panel, the robot transducer has now been developed using a 5-mm diameter, 2-mm thick PZT disc. This transducer has its lowest-order thickness resonance just under 1 MHz (at 960 kHz). The lowest-order radial resonance of the disc occurs at 460 kHz, and is reduced to ~ 400 kHz when bonded into the transducer housing.

The increased area of the PZT element drew attention to the importance of ensuring that its face sits flat on the aluminium surface. While the robot's foot has been

designed to try to assist this – by the inclusion of stops between each of the (flexible) suction pads onto which the applied vacuum should draw the foot – gravity, slight differences in the suction force at each pad, and surface irregularities on the aluminium would all contribute to preventing this happening. Consequently, to ensure maximum coupling, a flexible head for the transducer has been developed using a ball-and-socket coupling mechanism supported by a silicone rubber boot bonded to both ends of the inter-connecting parts. This head has been designed to give an angular variation of at most 5° in any direction with respect to the axis of the transducer (see Figure 4.8). Limiting this range is important to prevent any damage to the connections to the transducer's electrodes: the ball-and-socket construction and the flexible boot are designed to achieve this. This has proved to be a simple, inexpensive, and highly effective method of ensuring good coupling between the transducer face and the aluminium.

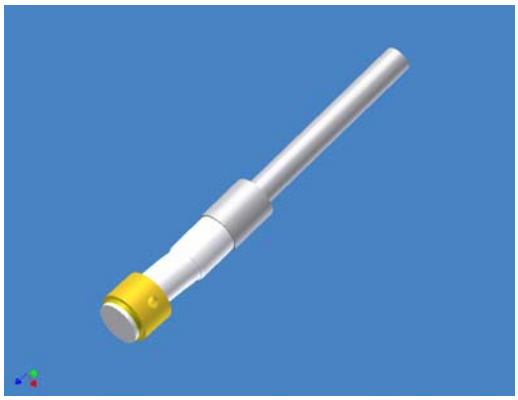

Figure 4.8: Robot transducer mounted on a flexible head.

Further improvements such as gold plating the front electrode and down the side walls of the transducer holder to effect the ground connection to the electrode have removed the possibility of having the thin copper ribbon connection (which previously acted as the earth line) break through constant flexing during operation of this transducer in the robot's foot.

4.3 Communication with the embedded agents

The 2.5-mm diameter, 0.5-mm thick PZT elements bonded to the inside of the aluminium panels are used as transmitters, both for communicating with the robot and, in principle, for active ultrasonic evaluation of damage. (A problem with the original design of the DAL boards prevents their use as ultrasonic receivers.) These elements have their lowest frequency resonance near 1 MHz, which is the fundamental radial mode of the disc. Their other major resonance is the first-order through-thickness resonance at about 4 MHz. In between these two are various higher order radial modes, but the separation of the first radial and through-thickness modes is sufficient to prevent significant mixing of these modes. The 1 MHz radial resonance, with a wavelength equal to the disc diameter (~ 2.5 mm), couples efficiently with the large in-plane component of the A_0 waveguide mode of the aluminium panel at this (λ, f) combination.

The electronics of the DAL board has low-pass filters on the receiver channels to attenuate any frequency components above 1.55 MHz. This is necessary because the analogue-to-digital converters on the Texas Instruments TMS320F2812 DSP chips are operated at a sampling frequency of 12.5 MHz (they have a maximum sampling rate of 16.666 MHz, but other constraints prevented operation at this rate). This allows each of the four channels (for each of the four PVDFs on each cell) to be sampled at 3.125 MHz, which means that the highest frequency that can be sampled without aliasing is about 1.56 MHz (the Nyquist frequency for this sampling rate).

For these reasons the A_0 guided mode of the aluminium panel at ~1 MHz, with a phase velocity of ~ 2.5 mm/μs and substantial in-plane and out-of-plane displacement components, was chosen as the communications channel.

The robot transducer has its lowest-order thickness resonance just under 1 MHz (at 960 kHz), and this couples well into the panel A_0 mode for communications. The lowest-order radial resonance of the robot's transducer occurs at ~ 400 kHz. This resonance also couples efficiently into the A_0 plate mode, which has a lower phase velocity (~ 1.75 mm/μs) at this frequency, and has been used for generating the signals that initiate a communications sequence and which enable the position of the robot foot to be determined by triangulation (see below).

The sequence of events for a communication between the robot and an embedded cell/agent is described in the following table.

Sequence	Embedded agent/cell action	Robot action
1	Agent awaits a detectable event.	Robot moves foot to cell location.
2		When robot foot is attached, the robot transmits a tone burst (5 cycles at 400 kHz) to initiate communications.
3	The agent detects a signal and decides whether it is a tone burst or an impact (see Section 3). (If an impact, it proceeds to impact location and diagnosis).	The robot switches its transducer electronics to receive mode, and waits to receive a packet of data.
4	Agent performs a triangulation procedure to determine the position of the robot transducer relative to the four PVDF sensors, based on the differences in the times of arrival of the tone burst.	
5	The agent transmits (using its central PZT transducer) a packet of data that contains: the coordinates of the robot transducer relative to the centre of the cell; and the gradient field values of the cell and of its connected neighbours.	

6	The agent waits for an acknowledgement signal.	The robot receives the data packet. If a packet is not received, it makes a small random move of the foot and tries again (2).
7		The robot sends an acknowledge signal and, based on the data in the packet, decides on its next action. Alternatives are as follows. Reposition its foot on the cell if it sufficiently far off-centre that it cannot make the next step. Step to the next cell as indicated by what it determines to be the highest priority gradient field values. Cease stepping and manipulate the camera (or other sensors).
8	If acknowledgement received, go back to (1). If not, wait for a pre-set time (~0.5 s), then re-transmit packet (5). If acknowledgement still not received, go back to (1).	

The times-of-arrival of the tone-burst signal at the agent's four PVDF sensors, which are required for triangulation as outlined in (4) above, are calculated by cross-correlation with a binary 400 kHz tone-burst filter. The peaks in the cross-correlations give the arrival times at the sensors. Triangulation is then done using a look-up table, as described in Report 5 [1].

The ultrasonic communications signals employ a 937.5 kHz carrier signal introduced into the aluminium skin. When the robot is transmitting, the carrier is generated by the robot's transducer and is received by the four PVDF impact sensors in the cell to which the relevant robot foot is attached. When the cell is transmitting, the carrier is generated by the cell's centre PZT transducer and received by the robot's transducer.

The carrier is Binary Phase Shift Key (BPSK) -modulated at a rate of 100 baud, which is slow enough to allow ultrasonic reflections within the panel to decay before the next symbol is sent. The effective data rate for the channel is approximately 80 bits/s, the reduced rate compared to the baud rate being due mainly to the synchronisation and checksum bits used.

4.4 Motion: stepping and navigation

The robot can move with six degrees of freedom: it can rotate each leg about a horizontal axis, and each foot has two independent rotations, one each about a horizontal axis and a vertical axis. The higher-level steps are defined in terms of these fundamental motions.

Because the robot has no global navigation capabilities and can only move from one cell to the next using dead-reckoning, large positional errors could rapidly accumulate as the robot moves over the surface. To avoid such positional errors, the underlying cell measures the robot's foot position by triangulation, as described above, and reports it to the robot. The robot can then either physically correct the foot position, or take it into account in calculating the step required for the next move.

A complication with this method of positional feedback is that the cell and the robot must have a common knowledge of the orientation of the cell relative to other cells and the structure. One way to avoid this issue is to build the CD structure with all cells in a prescribed orientation. However, it can be argued that this solution is inconsistent with the concept of an adaptive, self-organising structure, and a more satisfactory solution involves the cells cooperatively determining their relative orientations. This is also necessary for algorithms such as ant colony optimisation.

Nevertheless, there is a need for the robot to have some basic knowledge about the structure, since it cannot be allowed to step on the gaps between panels, and it needs to know where the face edges of the prism are located in order to be able to step from one face to another. A robot with more computational power than the present one could, for example, use its video camera to resolve local issues such as these, but for the time being the robot has been given this basic knowledge of the cell layouts.

The robot's navigation and functions will be determined cooperatively with the local agents embedded in the test-bed skin with which it is in contact. The robot will navigate around the surface of the test-bed using gradient field data available from the underlying cell to which it is attached at the time. This data is specific to the cell's local neighbourhood, and does not contain any global information about the system. Further information about the gradient fields is contained in Section 3 above.

4.5 Manipulation of camera and image data transmission

It was hoped that by the time of this report the robot would be using its video cameras to image the "damage" (at this stage a red or green spot stuck to the skin at the impact location), but the implementation of this has not yet been completed. The concept is that when the robot arrives at the neighbourhood of a damaged cell, as indicated by minimum in the gradient field, instead of putting its foot down on the damaged cell it will hold it above the cell in an orientation such that the video camera (Figure 2) can be directed at the damage. A routine has been written to recognise the colour of the damage indicator spot from its video image.

In the initial implementation of this damage imaging capability the data will be sent via a 2.4 GHz wireless link to the system visualiser PC for analysis, but the longer-term aim is for the image data analysis to be carried out on the robot. Use of the visualiser for this function would give it an essential role in the operation of the system, which is not intended. It also creates a potential single point of failure. While at present the robot represents a single point of system failure, the intention is to eventually have multiple (and large numbers of) robots that may share the damage evaluation task.

5. The System Visualiser

As described in earlier reports (e.g. [1], Sections 5 and 7), the system visualiser is a computer that is used for initialising the multi-agent system and displaying the state of the system, but which plays no essential role in the operation of the system. It can be attached via a USB port at a number of points around the edge of the CD system (in principle it could be anywhere), and its function is to request state information from the embedded agents and display this information.

The visualisation function has now been extended to show the robot position, and this is illustrated in Figure 5.1. This information is not obtained from the robot itself, which in principle doesn't need to know its absolute position on the structure, but from the agents with whom the robot is communicating. Thus, an observer doesn't need to be able to see the robot to be able to monitor its activities. This principle can of course be extended to more than one robot.

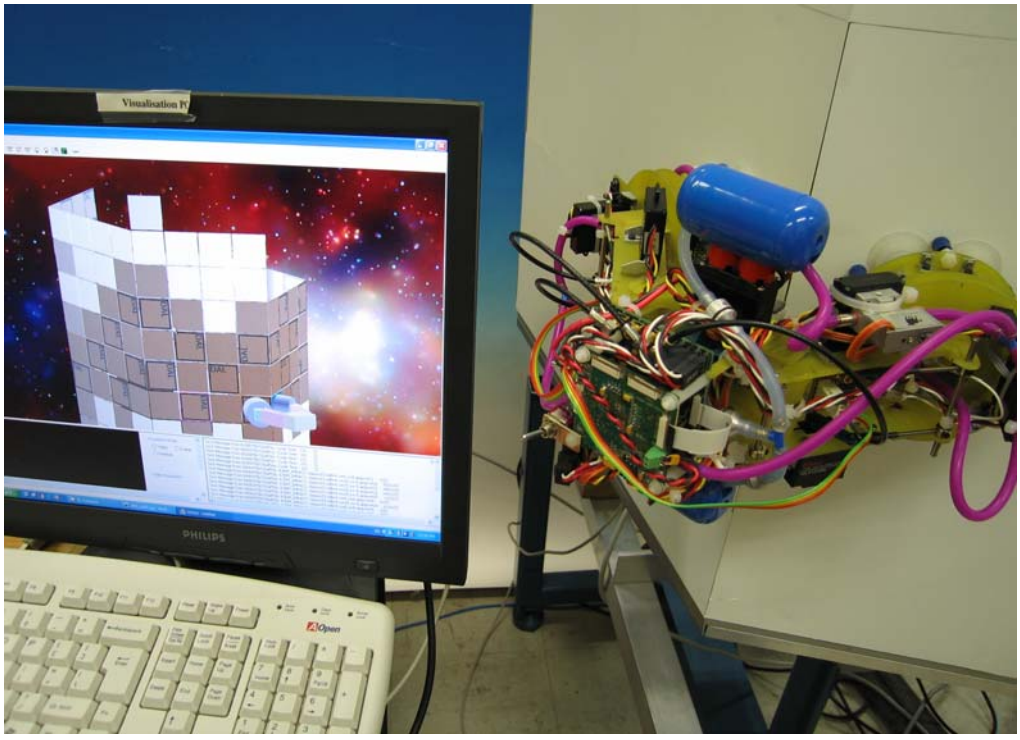

Figure 5.1: Visualiser screen (left), showing an animation of the robot in its actual position on the CD structure (right).

Four new views of the Concept Demonstrator have been developed to display and debug the gradient algorithm. Each view uses a different colour and represents the value of the gradient field on a particular cell by the shade of colour (see Figure 2.3). Lighter shades represent higher gradient values, which are further away from impact locations. Cells that the robot cannot reach because they have no DAL board attached are displayed in white. By double-clicking on a cell in the display the numerical values of the cell's gradients are displayed.

As was pointed out in the previous section, it is intended that, as an initial short-term measure, the visualiser will perform the role of analysing video image data from the robot to determine the severity of the damage. This is to be used to confirm (or otherwise) the damage diagnosis made by utilising impact data from the embedded piezoelectric sensors. The longer-term intention is for this analysis to be carried out by the robot (or robots).

This image analysis role is enabled by an image capture card with a 2.4 GHz wireless link to the video cameras on each foot of the robot. The images transmitted from the robot can be viewed in a new window that has been incorporated into the visualiser display. A Matlab routine that enables the colour of the "damage" indicator (as described in Section 3) to be identified has been compiled and added to the visualiser.

In the present simplified situation, identification of the colour of the damage indicator will provide an authoritative diagnosis of the damage that will supersede the tentative diagnosis made from data from the embedded piezoelectric sensors. However, the situation may not always be as clear cut as this. An optical image may underestimate internal structural damage within a material (e.g. in laminated composites), and the use of other sensor modalities for this secondary sensing may also give rise to ambiguities or uncertainties in some cases. In the longer term it will be desirable to make a diagnostic decision based on all of the available sensor data, possibly with the assistance of material or structural models. The use of data from multiple sensing modalities for damage diagnosis and prognosis will be the subject of future work.

In cases of suddenly occurring major damage, which will usually be the result of an external influence such as an impact, it may be necessary to initiate a response to an initial indicator (such as the sensing of the impact) without waiting for a subsequent detailed inspection and diagnosis. In such a perceived emergency situation a precautionary principle must clearly be adopted: rapid action should be taken first and a more detailed diagnosis made later.

6. Conclusions

This Report describes a significant advance in the capability of the CSIRO/NASA structural health monitoring concept demonstrator (CD). The main thrust of the work has been the development of a mobile robotic agent, and the hardware and software modifications and developments required to enable the demonstrator to operate as a single, self-organising, multi-agent system. This single-robot system is seen as the forerunner of a system in which larger numbers of small robots perform inspection and repair tasks cooperatively, by self-organisation.

While the goal of demonstrating self-organised damage diagnosis was not fully achieved in the time available, much of the work required for the final element – enabling the robot to point the video camera and transmit an image – has been either completed or planned. It is expected that this will be completed shortly. Nevertheless, what has been achieved is an important step. A demonstration video of the system operating will be made and forwarded to NASA.

The next steps in the development of the system are as follows.

- Completion of the present task of achieving camera operation and integration of diagnostic information.
- Adaptation of the system to incorporate a damage scenario of specific interest to NASA, namely monitoring the integrity of thermal protection tiles using a combination of acoustic emission and thermal evaluation. This work, which is being carried out in collaboration with NASA Dryden Flight Center, is now under way.
- Damage diagnosis based on data from multiple sensing modalities. This will be initially addressed through the work on thermal protection systems, but the development of a general framework for the use of multiple data sets for damage diagnosis is a significant problem.
- Expansion of the single-robot capability to accommodate multiple robots. There is a lot of current research interest in self-organised cooperating groups of robots for a wide range of applications in various environments.

While these are all substantial steps, the developments reported here represent an important advance and a sound base from which to make further progress.

7. References

1. A. Batten, 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, I. Sharp, D. A. Scott, P. Valencia, P. Wang, D. F. Whitnall (2004). *Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles. Report 5: Phase 2 Implementation of the Concept Demonstrator*, CSIRO Telecommunications and Industrial Physics. Confidential Report No. TIPP 2056.

2. D. Abbott, B. Doyle, J. B. Dunlop, A. J. Farmer, M. Hedley, J. Herrmann, G. C. James, M. E. Johnson, B. Joshi, G. T. Poulton, D. C. Price, M. Prokopenko, T. Reda, D. E. Rees, D. A. Scott, P. Valencia, D. Ward and J. G. Winter (2003). Concepts for an Integrated Vehicle Health Monitoring System, *Review of Progress in Quantitative Nondestructive Evaluation*, Vol. 22, pp.1606-14 (eds. D. O. Thompson and D. E. Chimenti), American Institute of Physics Conference Proceedings Vol. 657.

3. D. C. Price, D. A. Scott, G. C. Edwards, A. Batten, A. J. Farmer, M. Hedley, M. E. Johnson, C. J. Lewis, G. T. Poulton, M. Prokopenko, P. Valencia and P. Wang (2003). An Integrated Health Monitoring System for an Ageless Aerospace Vehicle, Structural Health Monitoring 2003, *Proceedings of the 4th International Workshop on Structural Health Monitoring*, Stanford, CA, September 2003. pp. 310-18. (ed. F-K. Chang) DEStech Publications.

4. J. Ferber (1999). *Multi-Agent Systems. An Introduction to Distributed Artificial Intelligence*. Addison-Wesley, London, UK.

5. D. C. Price, A. Batten, G. C. Edwards, A. J. D. Farmer, V. Gerasimov, M. Hedley, N. Hoschke, M. E. Johnson, C. J. Lewis, A. Murdoch, M. Prokopenko, D. A. Scott, P. Valencia and P. Wang (2004). Detection, Evaluation and Diagnosis of Impact Damage in a Complex Multi-Agent Structural Health Management System, *Proceedings of the 2nd Australasian Workshop on Structural Health Monitoring*, Melbourne, Australia. December 2004. pp. 16-27.

6. T. Kohonen (2001). *Self-organizing Maps*. Springer (3rd Edition).

7. T. Kohonen (2003). Self-organized Maps of Sensory Events, *Phil. Trans. Roy. Soc. Lond. A* Vol. 361, 1177-1186.

8. M. Prokopenko, P. Wang, D.A. Scott, V. Gerasimov, N. Hoschke, D.C. Price (2005). On Self-organising Diagnostics in Impact Sensing Networks *Proceedings of the 9th International Conference on Knowledge Based and Intelligent Information and Engineering Systems (KES-2005)*, Melbourne, Aust., September 2005.

9. M. Prokopenko, P. Wang, M. Foreman, P. Valencia, D. C. Price and G. T. Poulton (2005) On connectivity of reconfigurable impact networks in ageless aerospace vehicles. *Robotics and Autonomous Systems* 53, 36-58.

10. M. Hedley, M. E. Johnson, C.J. Lewis, D.A. Carpenter, H. Lovatt, D.C. Price (2003). Smart Sensor Network for Space Vehicle Monitoring, *Proceedings of the International Signal Processing Conference* (Dallas, Tx.), March 2003. http://www.gspx.com/GSPX/papers_online/papers_list.php

11. M. Dorigo and G. Di Caro (1999). Ant Colony Optimization: A New Meta-Heuristic, P*roc. 1999 Congress on Evolutionary Computation*, pp. 1470–1477, Washington DC, July 1999.

12. M. Dorigo and T. Stützle (2004). *Ant Colony Optimization*. MIT Press (Cambridge, Mass.).

13. D. Abbott, 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, D. F. Whitnall (2003). *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

Appendix A: Project Reports and Publications

This appendix lists the research output of the project group and collaborators since the beginning of the work in November 2001.

1. *Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles, Report 1*
   D. Abbott, S. Cunningham, G. Daniels, B. Doyle, J. Dunlop, D. Economou, A. Farmer, D. Farrant, C. Foley, B. Fox, M. Hedley, J. Herrmann, C. Jacka, G. James, M. Johnson, B. Martin, G. Poulton, D. Price, T. Reda, G. Rosolen, A. Scott, P. Valencia, D. Ward, J. Winter, A. Young
   CSIRO Telecommunications and Industrial Physics. Confidential Report No. TIPP 1516, 2001. Also published as NASA Technical Report no. NASA/CR-2002-211772, at http://techreports.larc.nasa.gov/ltrs/PDF/2002/cr/NASA-2002-cr211772.pdf

2. *Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles. Development of Concepts for an Intelligent Sensing System. Report 2*
   D. Abbott, B. Doyle, J. Dunlop, A. Farmer, M. Hedley, J. Herrmann, G. James, M. Johnson, B. Joshi, G. Poulton, D. Price, M. Prokopenko, T. Reda, D. Rees, A. Scott, P. Valencia, D. Ward, J. Winter
   CSIRO Telecommunications and Industrial Physics. Confidential Report No. TIPP 1517. May 2002. Also published as NASA technical report NASA/CR-2002-211773, at http://techreports.larc.nasa.gov/ltrs/PDF/2002/cr/NASA-2002-cr211773.pdf

3. *Concepts for an Integrated Vehicle Health Monitoring System*
   D. Abbott, B. Doyle, J.B. Dunlop, A.J. Farmer, M. Hedley, J. Herrmann, G.C. James, M.E. Johnson, B. Joshi, G.T. Poulton, D.C. Price, M. Prokopenko, T. Reda, D.E. Rees, D.A. Scott, P. Valencia, D. Ward and J.G. Winter
   Review of Progress in Quantitative Nondestructive Evaluation, Vol. 22, pp.1606-14 (eds. D.O. Thompson and D.E. Chimenti), American Institute of Physics Conference Proceedings Vol. 657, (2003).

4. *Experimental Observation of Linear and Non-Linear Guided Wave Propagation in Rolled Aluminum Sheets*
   D.A. Scott and D.C. Price
   Review of Progress in Quantitative Nondestructive Evaluation, Vol. 22, pp.1559-66 (eds. D.O. Thompson and D.E. Chimenti), American Institute of Physics Conference Proceedings Vol. 657, (2003).

5. *Theoretical Aspects of Linear and Non-Linear Guided Wave Propagation in Rolled Aluminum Sheets*
   D.C. Price and D.A. Scott
   Review of Progress in Quantitative Nondestructive Evaluation, Vol. 22, pp.1567-74 (eds. D.O. Thompson and D.E. Chimenti), American Institute of Physics Conference Proceedings Vol. 657, (2003).

6. *Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles. Report 3: Design of the Concept Demonstrator*
D. Abbott, J. Ables, A. Batten, D. C. Carpenter, A. F. Collings, B. Doyle, J. B. Dunlop, G. C. Edwards, A. J. Farmer, B. Gaffney, M. Hedley, P. Isaacs, M. E. Johnson, B. Joshi, C. J. Lewis, G. T. Poulton, D. C. Price, M. Prokopenko, T. Reda, D. E. Rees, D. A. Scott, S. Seneviratne, P. Valencia, P. Wang, D. F. Whitnall, J. Winter
CSIRO Telecommunications and Industrial Physics. Confidential Report No. TIPP 1628. January 2003.

7. *Smart Sensor Network for Space Vehicle Monitoring*
M. Hedley, M.E. Johnson, C.J. Lewis, D.A. Carpenter, H. Lovatt, D.C. Price
Proceedings of the International Signal Processing Conference (Dallas, Tx.), March 2003.

8. *Self-reconfigurable Sensor Networks in Ageless Aerospace Vehicles*
P. Wang, P. Valencia, M. Prokopenko, D.C. Price, G.T. Poulton
Proceedings of the 11th International Conference on Advanced Robotics (ICAR-03), Portugal, July 2003, pp.1098-1103.

9. *Self-organizing Impact Boundaries in Ageless Aerospace Vehicles*
H. Lovatt, G.T. Poulton, D.C. Price, M. Prokopenko, P. Valencia, P. Wang
Proceedings of the 2nd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2003), Melbourne, July 2003, pp.249-56 (ACM Press, NY).

10. *Phase Transitions in Self-organising Sensor Networks*
M. Foreman, M. Prokopenko, P. Wang
Proceedings of the 7th European Conference on Artificial Life (ECAL-03), pp. 781-791. Dortmund, Germany, September 2003.

11. *An Integrated Health Monitoring System for an Ageless Aerospace Vehicle*
D.C. Price, D.A. Scott, G.C. Edwards, A. Batten, A.J. Farmer, M. Hedley, M.E. Johnson, C.J. Lewis, G.T. Poulton, M. Prokopenko, P. Valencia, P. Wang
In "Structural Health Monitoring 2003", Proceedings of the 4th International Workshop on Structural Health Monitoring, Stanford, CA, September 2003 (ed. F-K. Chang) DEStech Publications, 2003, pp.310-18. (ISBN: 1-932078-20-7).

12. *Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles. Report 4: Phase 1 Implementation of the Concept Demonstrator*
D. Abbott, 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, D. F. Whitnall
CSIRO Telecommunications and Industrial Physics. Confidential Report No. TIPP 1898. Sept. 2003.

13. *Evolvable Recovery Membranes in Self-Monitoring Aerospace Vehicles*
P. Wang, M. Prokopenko
In Proceedings of the 8th International Conference on Simulation of Adaptive Behaviour (SAB-2004), Los Angeles, USA, July 2004, pp.509-18.

14. *On Connectivity of Reconfigurable Impact Networks in Ageless Aerospace Vehicles*
Mikhail Prokopenko, Peter Wang, Mark Foreman, Philip Valencia, Don Price, Geoff Poulton
Robotics and Autonomous Systems 53, 36-58 (2005).

15. *Self-organizing Hierarchies in Sensor and Communication Networks*
Mikhail Prokopenko, Peter Wang, Don Price, Philip Valencia, Mark Foreman, Anthony Farmer
Artificial Life (Special Issue on Dynamic Hierarchies) 11, 407-26 (2005).

16. *On Self-referential Shape Replication in Robust Aerospace Vehicles*
M. Prokopenko, P. Wang
Proceedings of the 9th International Conference on the Simulation and Synthesis of Living Systems (ALIFE9), Boston, USA, September 2004, pp.27-32.

17. *Testbed for Structural Health Monitoring Network*
M. Hedley, M. E. Johnson, C. J. Lewis, D. C. Price
Proc. Global Signal Processing Conference, Santa Clara, California, USA, Sept. 2004. ICT Centre Publication 04/1856.

18. *Communication Protocols for a Structural Health Monitoring Network*
M. Hedley
Proc. IEEE Conf. on Mobile Ad-hoc and Sensor Systems (MASS-2004), Fort Lauderdale, Florida, USA, October 2004.

19. *Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles. Report 5: Phase 2 Implementation of the Concept Demonstrator*
A. Batten, 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, I. Sharp, D. A. Scott, P. Valencia, P. Wang, D. F. Whitnall
CSIRO Telecommunications and Industrial Physics. Confidential Report No. TIPP 2056. April 2004.

20. *Intelligent Health Monitoring for Aerospace Vehicles. Report 1. Stage 1: Project Scope and Feasibility Demonstration*
I.S. Cole, D.C. Price, T. Muster, M. Hedley, R. Drogemuller, G. Trinidad, S. Egan, F. Boulaire, M. Prokopenko
CMIT, CTIP, CICT Centre Report to the Boeing Company. Confidential Report No. CMIT-2004-157, April 2004.

21. *An Intelligent Sensor System for Detection and Evaluation of Particle Impact Damage*

D.A. Scott, A. Batten, G.C. Edwards, A.J. Farmer, M. Hedley, N. Hoschke, P. Isaacs, M. Johnson, A. Murdoch, C. Lewis, D.C. Price, M. Prokopenko, P. Valencia, P. Wang
Review of Progress in Quantitative Nondestructive Evaluation, Vol. 24, pp.1825-32 (eds. D.O. Thompson and D.E. Chimenti), American Institute of Physics Conference Proceedings Vol. 760, (2005).

22. *Sensor Network for Structural Health Monitoring*
M. Hedley, N. Hoschke, M. Johnson, C. Lewis, A. Murdoch, D. Price, M. Prokopenko, A. Scott, P. Wang, A. Farmer
Proceedings of the 2004 Intelligent Sensors, Sensor Networks and Information Processing Conference (December 2004), pp.361-6 (2004). (ISBN: 0-7803-8894-1).

23. *Structural Health Management for Future Aerospace Vehicles*
W.H. Prosser, S.G. Allison, S.E. Woodard, R.A. Wincheski, E.G. Cooper, D.C. Price, M. Hedley, M. Prokopenko, D.A. Scott, A. Tessler, J.L. Spangler
Proceedings of the 2nd Australasian Workshop on Structural Health Monitoring (2AWSHM), Melbourne, Aust., December 2004 (ISBN 0-9756992-0-2).

24. *Detection, Evaluation and Diagnosis of Impact Damage in a Complex Multi-Agent Structural Health Management System*
D.C. Price, A. Batten, G.C. Edwards, A.J.D. Farmer, V. Gerasimov, M. Hedley, N. Hoschke, M.E. Johnson, C.J. Lewis, A. Murdoch, M. Prokopenko, D.A. Scott, P. Valencia, P. Wang
Proceedings of the 2nd Australasian Workshop on Structural Health Monitoring (2AWSHM), Melbourne, Aust., December 2004 (ISBN 0-9756992-0-2).

25. *Intelligent Health Monitoring for Aerospace Vehicles. Stage 2A Report*
I.S. Cole, P. Corrigan, D.C. Price, T. Muster, W. Ganther, M. Hedley, R. Drogemuller, D. Paterson, S. Egan, F. Boulaire
CMIT, CIP Report to the Boeing Company. Confidential Report No. CMIT-2004-523, December 2004. (75 pages).

26. *Development of a System for Corrosion Diagnostics and Prognostics*
Angela Trego, Don Price, Mark Hedley, Penny Corrigan, Ivan Cole, Tim Muster
Proceedings of 1st World Congress on Corrosion in the Military: Cost Reduction Strategies. (Sorrento, Italy, 6-8 June 2005).

27. *Complexity Metrics for Self-monitoring Impact Sensing Networks*
M. Prokopenko, P. Wang, D. Price
Proceedings of the 2005 NASA/DoD Conference on Evolvable Hardware, pp. 239-246. (July 2005) (ISBN: 0-7695-2399-4).

28. *On Convergence of Dynamic Cluster Formation in Multi-Agent Networks*
M. Prokopenko, P. Mahendra, P. Wang

Proceedings of the 8th European Conference on Artificial Life (ECAL-2005) (Canterbury, UK) Sept. 2005.

29. *On Decentralised Clustering in Self-monitoring Networks*
Piraveenan Mahendra rajah, Mikhail Prokopenko, Peter Wang, Don Price
Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems, pp. 1175-6 (2005) (ACM Press, N.Y.) (ISBN: 1-59593-093-0).

30. *Towards Adaptive Clustering in Self-monitoring Multi-agent Networks*
Piraveenan Mahendra rajah, Mikhail Prokopenko, Peter Wang, Don Price
Proceedings of the 9th International Conference on Knowledge Based and Intelligent Information and Engineering Systems (KES-2005) (Melbourne, Aust., Sept. 2005).

31. *On Self-organising Diagnostics in Impact Sensing Networks*
M. Prokopenko, P. Wang, D.A. Scott, V. Gerasimov, N. Hoschke, D.C. Price
Proceedings of the 9th International Conference on Knowledge Based and Intelligent Information and Engineering Systems (KES-2005) (Melbourne, Aust., Sept. 2005).

32. *Establishing a Physical Basis for the In-situ Monitoring of Airframe Corrosion using Intelligent Sensor Networks*
T. Muster, I. Cole, W. Ganther, D. Paterson, P. Corrigan, D. Price
Proceedings of the 2005 Tri-Service Corrosion Conference (Florida, USA, November 2005).

33. *Intelligent Health Monitoring for Aerospace Vehicles. Stage 2B Report*
I.S. Cole, F. Boulaire, P. Corrigan, R. Drogemuller, S. Egan, W. Ganther, T. Muster, D. Paterson, D.C. Price, M. Hedley, B. Hinton, S. Galea
CMIT, CIP, DSTO Report to the Boeing Company. Confidential Report No. CMIT-2005-520, December 2005. (137 pages).

34. *Self-Organising Impact Sensing Networks in Robust Aerospace Vehicles*
Mikhail Prokopenko, Geoff Poulton, Don Price, Peter Wang, Philip Valencia, Nigel Hoschke, Tony Farmer, Mark Hedley, Chris Lewis, Andrew Scott
In "Advances in Applied Artificial Intelligence", ed. John Fulcher (Idea Group, Hershey, PA, USA, 2006) pp.186-233.

35. *A Self-Organising Sensing System for Structural Health Management*
N. Hoschke, C.J. Lewis, D.C. Price, D.A. Scott, G.C. Edwards, A. Batten
Proceedings of the 10th International Conference on Knowledge-Based and Intelligent Information and Engineering Systems (KES-2006) (Bournemouth, UK, Oct. 2006).

Aug. 2006

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
01- 01 - 2010	Contractor Report	

4. TITLE AND SUBTITLE

Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles: Report 6 - Development and Demonstration of a Self-Organizing Diagnostic System for Structural Health Monitoring

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Batten, Adam; Edwards, Graeme; Gerasimov, Vadim; Hoschke, Nigel; Isaacs, Peter; Lewis, Chris; Moore, Richard; Oppolzer, Florien; Price, Don; Prokopenko, Mikhail; Scott, Andrew; Wang, Peter

5d. PROJECT NUMBER

NNL05AI81P

5e. TASK NUMBER

5f. WORK UNIT NUMBER

939904.05.07

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

NASA Langley Research Center
Hampton, VA 23681-2199

8. PERFORMING ORGANIZATION REPORT NUMBER

Report No. CIP 2544

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

National Aeronautics and Space Administration
Washington, DC 20546-0001

10. SPONSOR/MONITOR'S ACRONYM(S)

NASA

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

NASA/CR-2010-216186

12. DISTRIBUTION/AVAILABILITY STATEMENT

Unclassified - Unlimited
Subject Category 38
Availability: NASA CASI (443) 757-5802

13. SUPPLEMENTARY NOTES

Langley Technical Monitor: Edward R. Generazio

14. ABSTRACT

This report describes a significant advance in the capability of the CSIRO/NASA structural health monitoring Concept Demonstrator (CD). The main thrust of the work has been the development of a mobile robotic agent, and the hardware and software modifications and developments required to enable the demonstrator to operate as a single, self-organizing, multi-agent system. This single-robot system is seen as the forerunner of a system in which larger numbers of small robots perform inspection and repair tasks cooperatively, by self-organization. While the goal of demonstrating self-organized damage diagnosis was not fully achieved in the time available, much of the work required for the final element that enables the robot to point the video camera and transmit an image has been completed. A demonstration video of the CD and robotic systems operating will be made and forwarded to NASA.

15. SUBJECT TERMS

Nondestructive evaluation; Structural health monitoring; IVHM; NDE; NDI; NDT

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			STI Help Desk (email: help@sti.nasa.gov)
					19b. TELEPHONE NUMBER *(Include area code)*
U	U	U	UU	44	(443) 757-5802