Implementation Results for a Fault-Tolerant Multicellular Architecture Inspired by Endocrine Communication

Andrew J. Greensted and Andy M. Tyrrell
Intelligent Systems Research Group, Department of Electronics,
University of York, UK, YO10 5DD
tel:+44 1904 432379
http://www.bioinspired.com

E-mail: {ajg112, amt}@ohm.york.ac.uk

Abstract

The hybrid redundancy structure found at the cellular level of higher animals provides complex organism with the three key features of a reliability-engineered system: fault tolerance, detection and recovery. For this reason, both the operation and organisation of this redundancy scheme provide an attractive source of inspiration for an electronic fault tolerant system.

The electronic architecture documented within this paper models the cooperative operation and consequent fault masking of the multiple cells that form biological organs. A communication system, inspired by endocrinology, is then used to network together these cells, coordinating their activity as organs, and controlling the operation of data processing tasks on a data stream.

The bioNode hardware platform is used to implement and test the presented endocrinology inspired architecture. Results of the system's operation are provided to demonstrate the architecture's ability to maintain correct computation on a data stream whilst being subjected to multiple and varied hardware faults.

1. Introduction

The examples of redundancy found in biological organisms are both extensive and varied. As a source of inspiration, the different methods of redundancy organisation and utilisation provide an attractive pool of techniques and architectures applicable to the development of novel reliability engineered electronic systems. Of particular interest are cases where the inclusion of redundancy promotes fault tolerance in biological systems. It can be observed that this characteristic is present throughout the various levels of biology's structural hierarchy.

For example, at a low level, genetic redundancy, in the form of tandem arrays found in DNA, consists of multiple genes each specifying a single whole protein. Genetic faults are then able to occur without loss of the protein's expression [10, 1]. At a higher, social, level in the hierarchy, ant foraging systems can suffer individual loss without stopping the flow of food, and thus failure of the system as a whole [4]. With such a diversity of fault tolerant biological systems, the question that arises is which may be the most appropriate for inspiring an electronic hardware architecture?

This paper presents the design and operation of a novel multicellular fault tolerant electronic architecture, the inspiration for which is drawn from the operation and organisation of the cellular level of higher-animals. Cellular redundancy provides both structural and functional robustness to their host organism. Accordingly, minor loss of cells neither diminishes the structural integrity of elements such as blood vessels and connective tissues, nor does it negatively impact the operation of vital organ functions. It is this latter point of active functional redundancy that is a key motivation for the presented research.

Taking a simplified viewpoint, the many cells that make up a single organ operate in a concurrent and coordinated fashion when performing the organ's specific function. If that organ is considered to have a homogeneous construction of functionally identical cells, then a fault can occur in any of these constituent cells without loss of overall function. A simplified implementation of this redundancy architecture forms the basis for the novel electronic architecture in which multiple artificial cells are formed into organ groups.

In biological systems, cells and organs do not operate in isolation. The coordination of cells within organs, and the communication between organs is a necessary part of creating a full living organism. Therefore, simply recreating an artificial cellular structure hardly constitutes a functional

and useful system. To this end, further biological inspiration has been used to provide an overlaying communication system that utilises the underlying cellular architecture.

Endocrinology is the study of chemical messengers, such as hormones, and the glands that secrete them into the body [3, 6]. The chemical messengers mediate a number of inter-cell communication techniques, two of which are endocrine and paracrine communication. These communication methods provide the means to link together the processing operations of the architecture's organs and achieve the required inter-cell coordination.

The overall result is a simplified biological system that can map into a hardware design; an underlying fault tolerant cellular architecture and an overlaying communication system, which together create a data stream processing system.

Section 2 documents the mapping of these biological systems into electronic hardware; mainly the creation of the simplified cellular architecture and the incorporation of endocrine and paracrine communication. Section 3 describes the physical implementation of the resulting design. Section 4 demonstrates the system's operation including an analysis of its fault tolerant behaviour. Concluding remarks are made in Section 5, and suggestions for furthering this research in Section 6.

2. The Architecture

The architecture presented within this paper can be considered as two layers, and shall be presented as such. Section 2.1 describes the underlying cellular architecture inspired by biological redundancy found at the cellular level. Section 2.2 describes the overlaying communication system, inspired by endocrinology, that converts the inert cells into a data processing system.

2.1. The Underlying Cellular Lattice

Approximations for the number of cells in the human body vary, but is thought to be in the region of 50×10^{12} to 100×10^{12} [2, 14]. Accordingly, the number of cells involved in the construction of animal organs is of considerable magnitude. Even with current electronic fabrication techniques, capable of placing many millions of transistors into a single device, reproducing this quantity of processing units of equivalent complexity to cells is unrealistic. Therefore a much simplified cellular structure is required, such as that shown in Figure 1.

Despite the relative simplicity of this architecture, it still incorporates the core features necessary to recreate the fault tolerance of the biological system. Artificial organs are constructed from groups of cells, all specialised to perform their organ's required operation. By spatially grouping together the cells of each organ, as would be the case in biology, cells are able to coordinate their operations as a whole. Subsequently, in the event of a biological fault in reality, or a hardware fault in the artificial case, the remaining cells within the group will maintain the organ's operation.

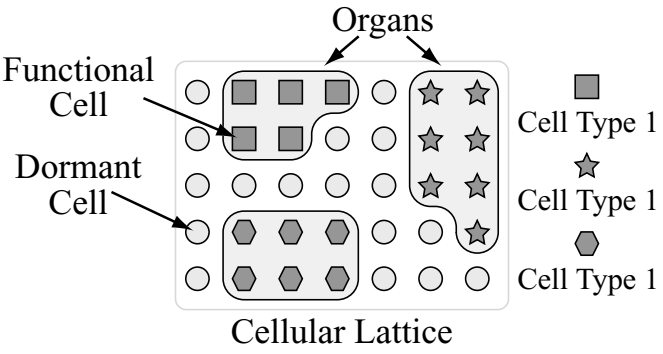

Figure 1: The simplified cellular lattice.

In either case, this is in essence a static redundancy scheme with similarities to an NMR (N-modular redundancy) architecture [11]. Of the many cells performing in parallel, the majority action far outweighs any minority of failed cells that act in contrast. The benefit of the biological system is the flexibility in which cells can be placed within the voting group and replaced if found faulty. In fact, the redundancy arrangement of cells in biological organs more closely operates in a hybrid form, due to the body's ability to detect faults using the immune system [9] and then recover from them by cell regrowth [17]. Unfortunately, current electronic fabrication is again unable to offer a viable artificial alternative to regrowth, and so circuitry replacement must be accomplished by the incorporation of spare hardware within the architecture.

As shown in Figure 1, the basic cellular lattice includes a number of dormant cells. These act as the spare units that may be incorporated into an organ that has suffered cell loss, replenishing its complement of operational cells. Rather than limiting a cell's potential to join any organ by pre-specialising it to a particular organ operation, dormant cells are capable of performing all the tasks required by any organ in the system. The biological inspiration for this feature is pluripotent stem cells [5]. These have the ability to specialise into any cell in the body, a behaviour made possible by all cells carrying a complete organism genome. This in itself is not a new technique, and has been used in architectures such as embryonic arrays [13, 16].

2.2. The Overlaying Endocrinology Control

To survive, our bodies must adjust to all the physiological changes inflicted upon them. Changes in environmental factors, such as temperature and humidity, require responses to maintain the equilibrium of our physiology. The primary role of the endocrine system is to provide this maintenance system. In order to achieve homoeostasis, the operation and state of cells must be controlled. This is achieved using a number of inter-cell communication techniques.

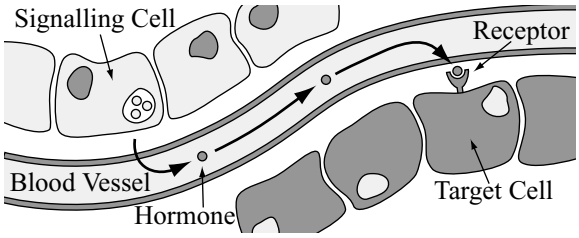

Figure 2: Endocrine Signalling. A long distance public form of inter-cell communication mediated by hormones.

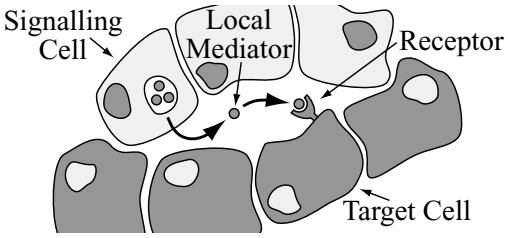

Figure 3: Paracrine Signalling. A localised form of inter-cell communication mediated by local messengers.

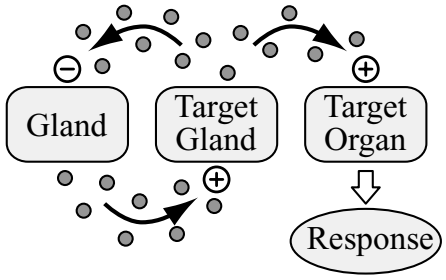

Figure 4: A basic endocrine control scenario. Single hormone types can invoke both stimulatory and inhibitory responses depending on the target reached.

An example of endocrine control in operation is the regulation of blood glucose levels. Regulation within a narrow allowable band is controlled by two hormones, insulin and glucagon. To decrease the level of glucose, the messenger insulin is released from the pancreas. This messenger binds to cells, causing them to absorb glucose through special channels on their surface. To increase the glucose level, glucagon is released from the pancreas which stimulates the liver to produce glucose.

In both cases endocrine communication is used to allow one group of cells to control another. Figure 2 depicts this communication process. Of all the inter-cell communication techniques available to a multicellular organ, endocrine signalling is by far the most public and generally the longest travelling. The source cells release their messenger signals, in this case called hormones, into the blood stream, in which they can travel to all parts of the body to find their target. The correct target cells are selected via a lock and key type binding mechanism, where the messengers have the correct key-like formation to fit the lock of the target cell's receptor.

Whereas endocrine signalling allows organs to communicate across the body, paracrine signalling operates over much shorter distances. This provides a means for local inter-cell coordination. As shown in Figure 3, paracrine signalling operates in very much the same way as endocrine signalling. However, the signalling molecules, in this case called local mediators, do not enter the blood stream. To maintain the locality of paracrine signals the signalling molecules are very short lived. If they are not rapidly taken up by their target cells they may be destroyed by enzymes in the extracellular space. Movement may be further restricted by the extracellular matrix, the physical structure around the cell.

2.2.1 Utilising Endocrine Control

The task of homoeostasis control generally requires using endocrine communication within a set of feedback paths, an example of which is shown in Figure 4. The key features of this biological process are the use of multiple signalling molecules to control the state of cells in different organs and the ability of the same signalling molecule type to both stimulate and inhibit a reaction. In the case of Figure 4, the initiating gland stimulates a target gland with the first hormone type. This first target gland then releases a second hormone type which inhibits the initial hormone release but also stimulates the final target organ, achieving the required response.

This very basic control scenario can be artificially extended to form a chain of successively activated organs, each in turn releasing specific hormones to activate the next, and inhibit the previous organ. An illustration of this system is shown in Figure 5a. The conceptual leap that converts this extended system into a useful electronic platform is the assignment of a computational operation to the cells of each organ (Figure 5b). The incorporation of data into the inter-organ messages allows information to be passed along the organ chain as arguments for the successive operations.

Using endocrine communication to bind together the sequence of organ operations has a number of fault tolerance advantages. The lack of hardwired communication links between source and target cells means the signalling molecules that mediate the messages have freedom to change their path, allowing for circumvention of problem areas within the cellular lattice. Furthermore, the use of multiple messenger molecules in a single communication means low level loss or corruption of molecules will not automatically result in overall communication failure. Secondly, the feedback mechanisms employed during endocrine control provides an acknowledge signal from the target to source organ. In practise, the source organ will

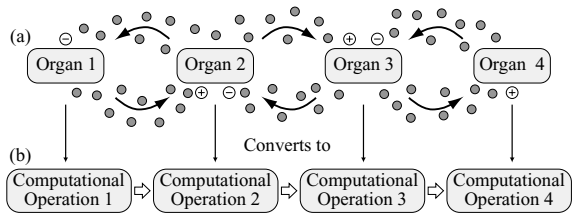

Figure 5: An artificially extended endocrine control cascade. The assignment of computational processes to organs creates a data processing pipeline.

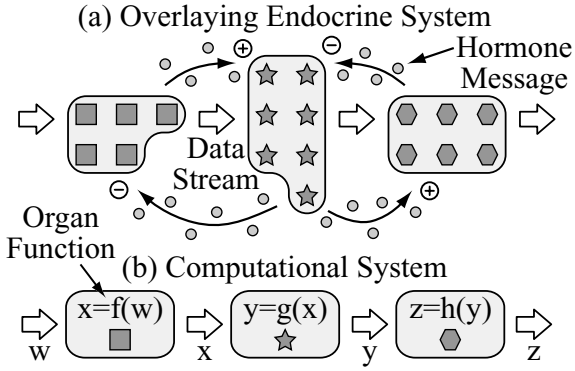

Figure 6: The final architecture. Cell groups performing specific computational tasks pass data using endocrine signalling.

continue to release signalling molecules until it receives the acknowledgement that the target has been reached, making for robust communication. Finally, endocrine communication makes good use of the underlying cellular redundancy. The use of multiple source cells to create the signalling messages, and similarly, multiple cells as targets means inter-organ communication, and therefore the flow and processing of data, can continue with the presence of cell hardware faults. What is more, the freedom of signalling molecule movement reduces cell placement restrictions, allowing any target location to be reached.

A summary of the developed architecture's structure and operation is depicted in Figure 6. Clusters of electronic cells positioned on the simplified cellular lattice form organs. Using electronic messages that travel through the lattice, organ groups can communicate in an endocrine signalling fashion. In doing so they are able to pass data through the system to be operated upon by a succession of activated organs. The process and design decisions necessary to convert this into an operating hardware system are presented next in Section 3.

3. The BioNode System

Preliminary simulations of the architecture developed in Section 2 positively indicated the system's correct operation [7]. However, a full hardware implementation can offer a better platform for testing and verification. For this task a multiprocessor hardware platform was created capable of implementing the architecture.

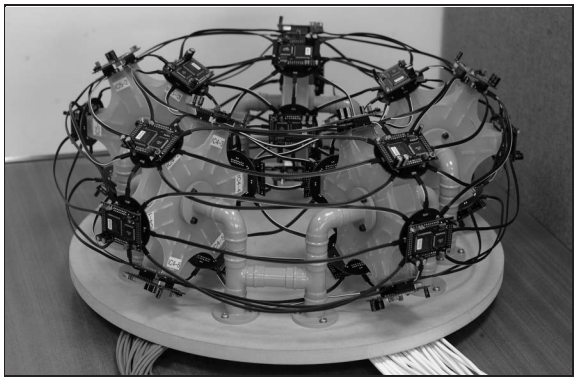

Figure 7: The bioNode System. A nearest neighbour network of thirty microcontrollers.

The bioNode system comprises a network of highly reconfigurable data processing modules combined with a set of supporting sub-systems, all custom built at the University of York (Figure 7). Each individual bioNode module consists of an Atmel™ microcontroller and a Xilinx™ FPGA, providing scope for both hardware and software reconfiguration [8]. The system was designed specifically for implementing multi-node bio-inspired systems, and therefore provides a perfect platform for testing this cellular biology and endocrinology inspired architecture.

3.1. An Artificial Signalling Space

The square grid upon which the cellular architecture is based (Figure 1) must be able to support an electronic version of endocrine signalling. Therefore, a method of transferring signalling molecules, or data packets, from cell to cell is required. To simulate this message space the inter-bioNode connections are used. Each bioNode has a connection to its eight nearest neighbours with which it manages a virtual signalling space around itself. Messages, analogous of signalling molecules, are able to pass from space to space via the communication links. Messages within a bioNode's space are subject to binding to the host bioNode in a manner similar to molecule-receptor binding. Messages contain, along with their data payload, a binding signature which is set by the cell that initially created the message. On release, messages travel around the lattice until they are received by a bioNode that is sensitive to the message's binding signature, at which point the target is found, and the message consumed.

To ensure messages are able to reach different bioNodes in the lattice, an artificial blood flow is introduced into the message passing system. BioNodes use a weighted random decision to choose which neighbour should receive a message. The weighting is universal to all bioNodes and is biased to one direction of transmission. The result is a

net direction of message flow much like the circulation of blood that performs the same task in biology. The message circulation is improved by removing the boundaries from the cellular lattice, which is why the bioNode network has a toroidal shape.

As in biology, messengers must have a finite lifetime to avoid their continual circulation in the event they repeatedly do not bind. In biology, signalling molecules naturally decay or are actively removed from the body. An artificial method of ageing messages is required, but without a central point of control, a time reference common to all cells is not available. Instead, each message stores a count of the hops it has made between bioNodes. When the count exceeds a limit, the message is deemed to have expired and is removed.

3.2. Further Implementation Issues

The following points describe further implementation issues:

• As mentioned in Section 2.1 the number of cells used when recreating a cellular lattice in electronic hardware has to be greatly reduced. This poses the problem that the masking effects are significantly reduced due to the potential lesser ratio of healthy to failed cells. To counter this problem a majority voting stage is introduced to the cell binding process. For a bioNode to become activated it must receive stimulatory messages from three unique sources. Similarly, to become deactivated, three inhibitory messages of diverse source are required. This improves the fault masking effect as messages released by failed cells will not be matched and will be ignored. A unique identification number stored by all bioNodes is tagged to a message's binding signature when created. It is these values that are used by subsequent bioNodes to distinguish between different message sources.

• Cell deactivation normally occurs when adequate inhibitory messages reach the organ. However, being the last organ in the chain, the cells of the output organ have no source of inhibitory signals. They therefore require an alternate method to become deactivated so that they may become available for operation on the next data sample. Automated deactivation is achieved by assigning a cell with a level of activation when it becomes stimulated. This activation level decays with time and if allowed to reach zero before signal induced inhibition occurs, will deactivate the cell. As well as solving the output organ problem it also stops any cells from remaining activated if they fail to receive the required inhibition messages.

• Inputting the data stream into the system involves passing data samples directly to the cells of the input organ and placing them into immediate activation. Output data is retrieved from the cells of the last organ, and combined to form a final output value. This initial input and final output of data is achieved using a set of communication cables, one for each bioNode, connected to an external system. Apart from this task, these links are mainly used for monitoring the internal status and operation of each bioNode, and not as a central point of control.

• Coordinating the operation of all an organ's cells is a necessary task if there is to be consistency in the sequence of data samples operated upon. If some cells are activated to operate on a data sample whilst others are not, the sequence of data seen by each cell will differ and may even be incomplete. For implementing advanced computational systems, reception of a consistent succession of data samples is vital. The introduction of paracrine signalling can be used to achieve cell coordination across the organ such that all cells become activated and operate on the same data sample together. On the binding of a message, the bioNode sends a 'single hop' copy of the message to its neighbours. The spatial locality of an organ's cells improves the chances that this message will bind to a neighbour and invoke a reaction.

• To create a cell's ability for pluripotence it must be able to perform all the computational operations required by any organ in the system. The use of a microcontroller for performing cell operations makes this task simple. Once a cell is assigned to an organ it simply needs to call the relevant software routines to perform the required operation. Furthermore, this approach of using software to implement cell function greatly increases the functional flexibility of the architecture.

3.3. Fault Injection

To test the architecture's ability to tolerate faults, a number of fault injection methods have been included into the bioNode system. These may be categorised into those that may be manually introduced and those that are injected programmatically. Manually creating faults can be achieved due to the pluggable nature of the bioNode system. Firstly, the interconnection wires used for inter-bioNode communication are removable and thus can induce communication failures. Secondly, each bioNode is connected to the system network via a passive connection board allowing the bioNodes themselves to be removed from the network. As well as this being a method of fault injection, it allows bioNodes with real faults to be replaced.

The second category requires the transmission of fault injection commands to the target bioNode from an external system. The transmission of these commands, over the bioNode's monitor link, is controllable via software and can be automated. The type of faults that may be injected are described in Section 4.2 when the architecture's fault tolerance abilities are demonstrated.

4. Experiments and Results

The results of two experiments are included in this section to demonstrate two key points of the developed

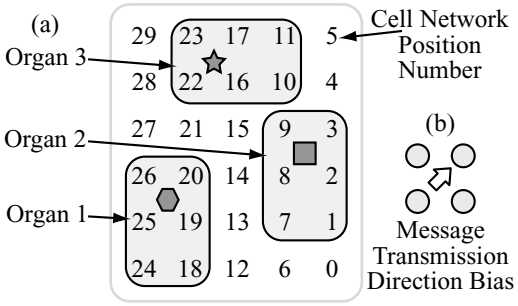

Figure 8: The cellular organisation of the test system. Three organs, each with six cells.

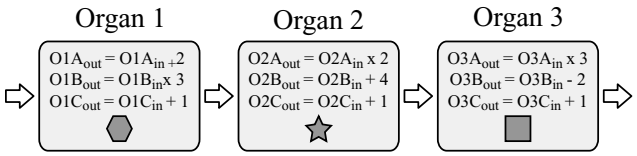

Figure 9: The computational pipeline. Organs perform three separate operations on different datums held within a single data stream.

architecture's operation. First is a demonstration that the required cascade of organ activation is achieved, thus achieving the correct series of computation on the data stream. The second set of results serves to demonstrate the system's ability to tolerate faults.

The cellular organisation used for these tests is depicted in Figure 8. The system's overall computational task is divided among three organs, each containing six cells. The data samples that construct the data stream each contain three datums. Each cell has a specific computational operation per datum, shown in Figure 9, creating a set of sub-streams labelled A to C.

4.1. Cascade Operation

The first experiment demonstrates the system's basic operation. Figure 11 shows the length and level of cell activation with time. Figure 10 provides a key for this plot. It can be observed that the cells of organ one, network positions 18, 19, 20, 24, 25 and 26, always reach activation at the same time. This is due to them being stimulated by the external data source when new samples are input to the system. It is also possible to see the process of cell deactivation by inhibition signals. The activation time of organ one and two cells is normally shorter than that required for automatic deactivation to occur. This is not the case for organ three cells, that require auto deactivation due to their lack of incoming inhibition signals.

Despite the use of paracrine signalling to ensure cell coordination, there are still occasions when cells miss activation for a data sample, such as Cell 10, data sample 7. For some forms of computational operations this could

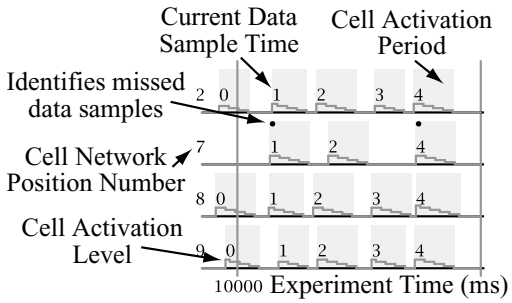

Figure 10: Key to Activation Plot. Timing variations can occur due to differences in status data arrival times between BioNodes.

invalidate the cell's calculated result and therefore must be minimised. This is a topic for further work and is discussed in Section 6.1.

4.2. Fault Tolerance

In this second experiment the architecture's fault tolerant abilities are tested by injecting faults into a number of bioNodes. The faults used are:

- *Complete bioNode shut down*. The bioNode enters an inactive state in which it neither releases nor receives messages.
- *Garbage Message Data*. The bioNode is forced to output messages of correct packet structure, but carrying random information.
- *Prolonged bioNode Activation*. Once activated the bioNode ignores all inhibitory messages and uses a longer auto deactivation time.

Figure 12 shows the activation levels of the cells during this experiment. The labels along the top of the plot indicate when, what type and to which bioNode a fault is injected. Figure 12 only shows the activation levels of those cells assigned to an organ. Although dormant cells do not become activated, they do participate in the system's operation. To maintain the free flow of messages, dormant cells must also receive and retransmit messages, and therefore they are candidates for fault injection. At experiment time $T = 17860ms$ the dormant bioNode at network location 21 is shutdown, without any effect on overall system operation.

This experiment also shows the architecture's ability to reintegrate cells into an organ. The bioNode at network location 8 is shutdown after data sample 2. It is restarted at data sample 17 and can be seen to resume normal service. The same process occurs when incorporating a cell from a different network location into an organ. After a faulty cell is shutdown, a dormant cell can be successfully incorporated into an organ, demonstrating the architecture's ability for reconfiguration.

To demonstrate that during the injection of these faults the system continues to operate, Figure 13 shows the system

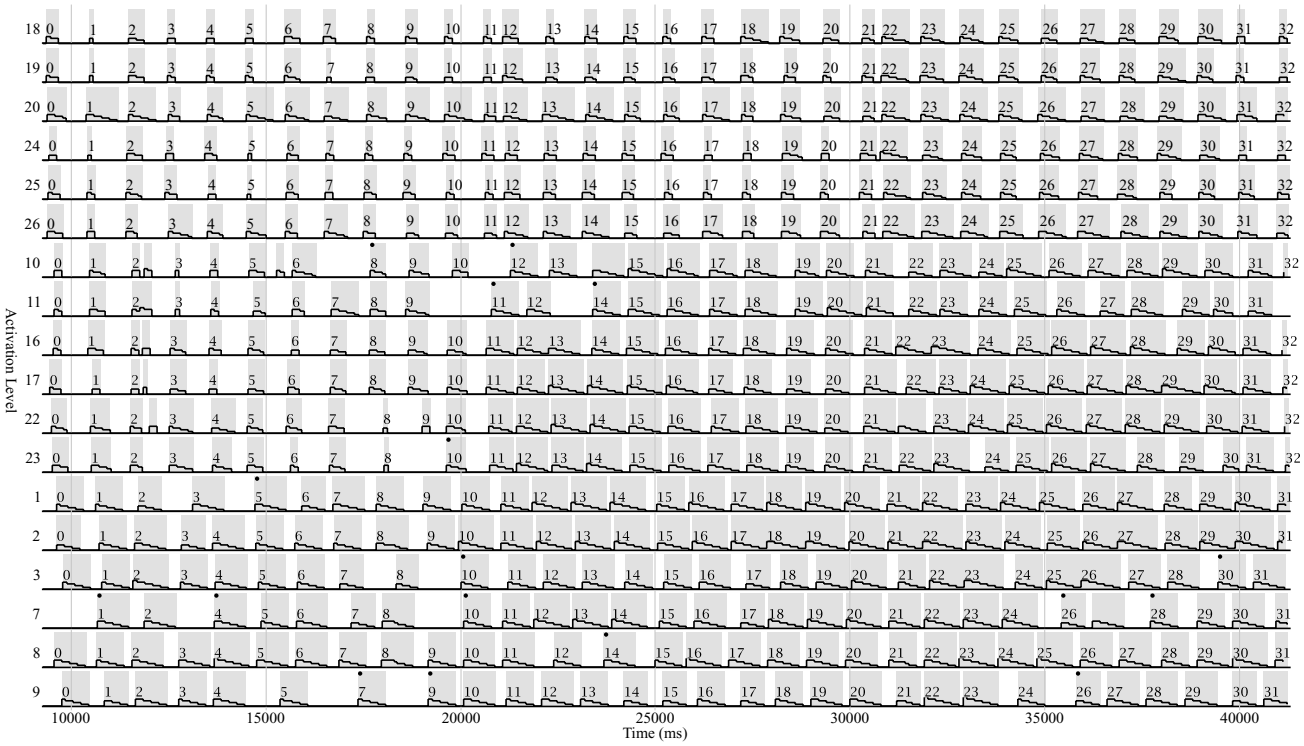

Figure 11: Experiment 1. Cell activation levels in a fault free system

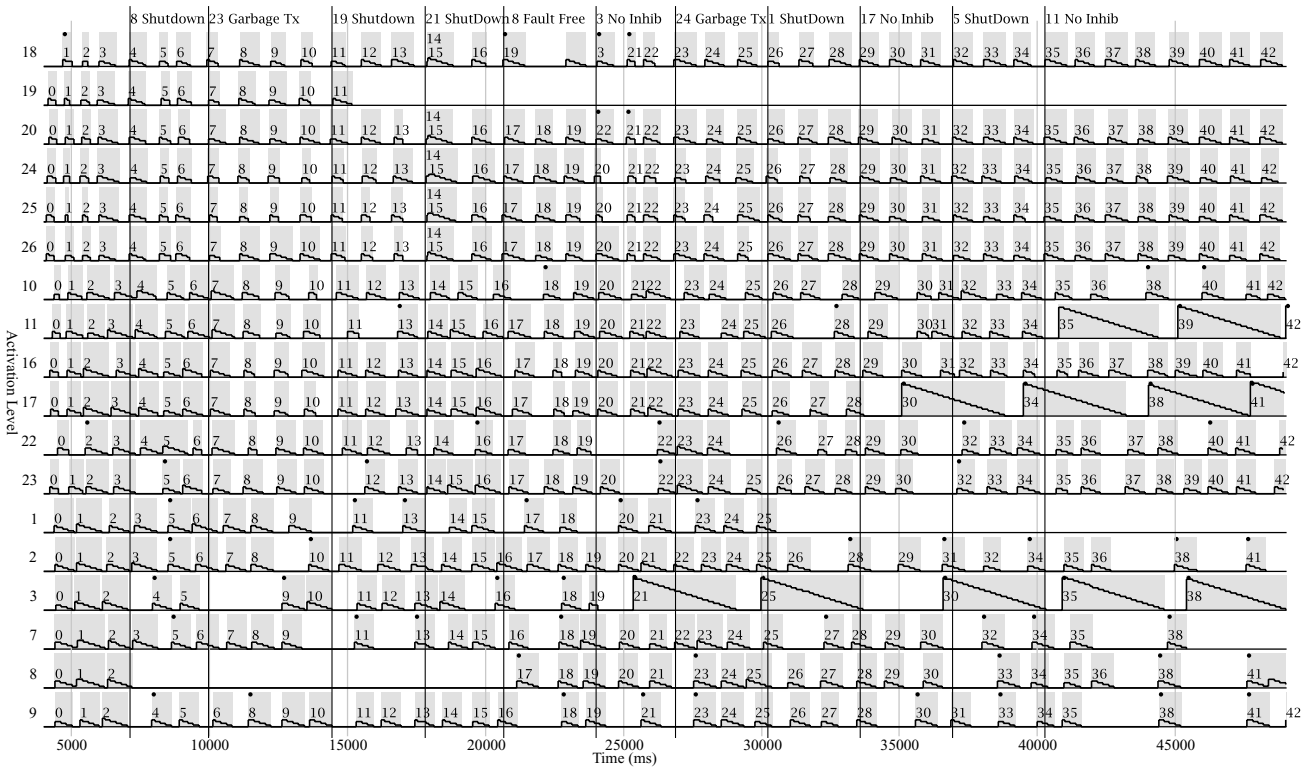

Figure 12: Experiment 2. Cell activation levels, with error injection

Each horizontal plot shows the changes in activation level for a single bioNode. The length of activation, per data sample, is shown by the grey shaded areas. Black dots in some of these areas represent a skip in processed data samples. See Figure 10 for a general key.

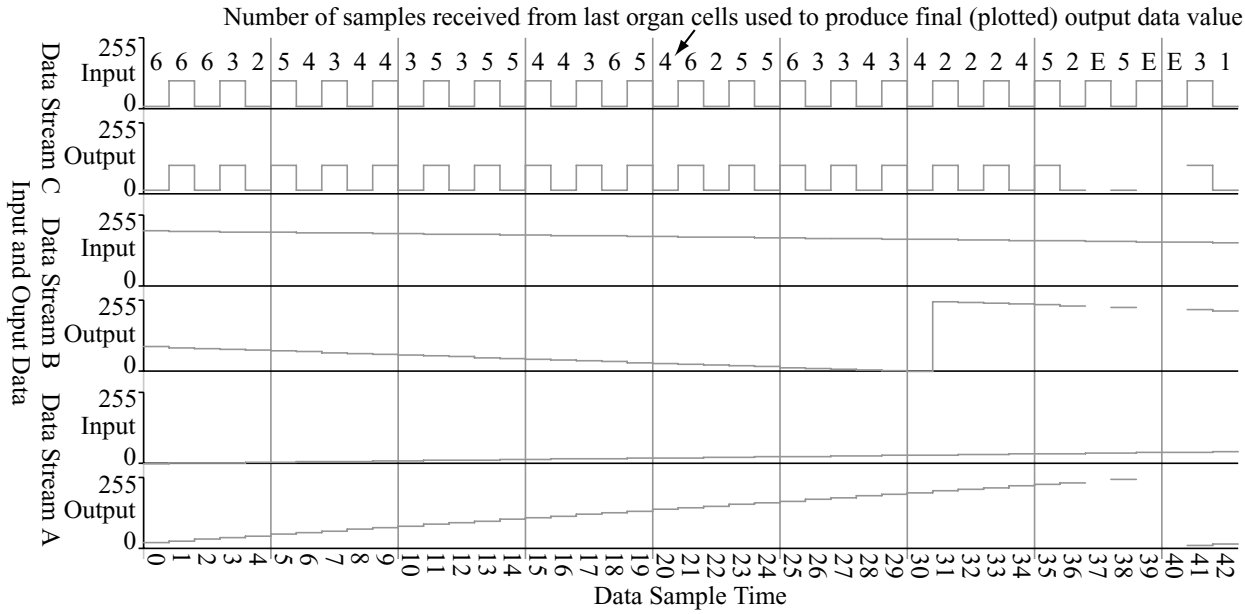

Figure 13: Experiment 2. System data input and output, with fault injection

input and output data plotted against the input data sample number. The system is able to output a computationally correct data stream until sample 37, after the injection of 10 faults. After this time, the external system that receives the output data from the cells of the last organ is no longer able to consistently produce a valid majority vote on the data it has received.

5. Summary and Conclusion

The ability of natural systems to survive injury and environmental change has always been a major source of inspiration in engineering, especially in terms of reliability and fault tolerance. The architecture demonstrated within this paper recreates a small portion of biological fault tolerance tailored to perform useful computation. The overall aim is to consider what low-level structural features of biology can create a fault tolerant foundation upon which a higher-level biological system can operate, taking advantage of the underlying architecture and producing usable functionality.

Three different sources of biological inspiration form the final architecture. Cellular redundancy inspires the platform's hardware organisation, the static redundancy scheme of multiple cells operating as an organ producing fault masking. Endocrinology inspired communication methods link together each organ's computational operations into a useful and fault tolerant system. Finally, cellular pluripotence enables dormant cells to be incorporated into any organ type that may have a deficit of functional cells.

The implementation and operational results of this architecture using the bioNode system demonstrate that a computational system can be created in the described manner. Furthermore, the system is able to tolerate varied and multiple hardware faults without overall system failure and has the capacity to recover from them.

A number of biology to electronics conversion problems have been tackled and solved. Mainly, dealing with the reduced redundancy when recreating a cellular structure, the recreation of endocrine feedback control and the utilisation of paracrine signalling for cell coordination.

Development of this architecture does requires further work to discover its full potential as a computational system. The results so far have provided a promising start, indicating the merit of this form of biological inspiration.

6. Further Work

This section outlines directions for further research.

6.1. Improved Cell Coordination

The coordinated operation of cells is of vital importance if organs are to operate as a whole, performing their computational task on every data sample that passes through the system. As mentioned previously, ensuring this coordination is important if the architecture is to implement calculations that operate on a history of consecutive samples. The incorporation of paracrine communication makes a significant improvement to coordinating cell activation, however further measures are required to reduce as far as possible the chance of skipped data samples during normal operations.

The problem is rooted in the timely arrival of enough messages of disparate source to cause cell activation. There

are two evident solutions. The first is to reduce the number of different cell sources required to activate a cell. If the number is reduced to two, activation may occur more readily, but there is more chance for incorrect activation due to failed source cells. Secondly, the effect of paracrine messages can be improved so that a successful binding of signal to cell has more affect on the neighbouring cells. A likely side effect of this approach is an increase in network traffic. Both approaches are yet to be tested.

6.2. More Complex System Applications

If a successful implementation of improved cell co-ordination can be achieved, applications of far greater complexity may be implemented. Applications that require a history of data samples include filters, integrators and differentiators. This opens up the usefulness of the proposed architecture to perform a multitude of DSP and automated control tasks.

In terms of hardware utilisation of the bioNode system, more complex computational operations make much better use of each bioNode's processing power; thus lessening any imbalance between communication overhead and useful computation.

6.3. Automated Development and Fault Detection

Currently the assignment of cells into organs and the spatial placement of organs on the cellular lattice is performed manually. In biology, this type of process is taken care of during the organism's development. As an embryo grows, environmental factors and a cell's position determine its path of specialisation. Using an artificial form of this process for automated assignment of cell operation within the cellular lattice would be an interesting direction for further work, and could perhaps follow the work of [12, 15]. A developmental process that could automate the inclusion of dormant cells into deficient organs would further enhance the fault tolerant abilities of the architecture by converting the otherwise static scheme into a hybrid redundancy scheme.

Automating the cell replenishment process requires a fault detection mechanism. Cells would require a method of identifying neighbours that have failed, invoking a form of apoptosis upon it and finally switching in a replacement of the same type. Detection of failure could be achieved by monitoring message data. Receiving cells can check the integrity of a incoming message and detect if the transmitting cell is producing erroneous data. As all cells have the ability to perform any system operation, it may be possible for cells to check the computational result stored in a message.

Once a cell has been detected as failed, it can be isolated from the network by ignoring all messages incoming from, and stopping all outgoing messages to the failed cell.

7. Acknowledgements

This research was undertaken with sponsorship from the EPSRC, grant reference: GR/R74512/0

References

[1] B. Alberts, A. Johnson, J. Lewis, M. Raff, K. Roberts, and P. Walter. *Molecular Biology of the Cell*. Garland Science, 4th edition, 2002.

[2] C. Blakemore and S. Jennett, editors. *The Oxford Companion to The Body*. Oxford University Press, 2001.

[3] C. G. D. Brook and N. J. Marshall. *Essential Endocrinology*. Blackwell Science, 4th edition, 2001.

[4] S. Camazine, J.-L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, 2001.

[5] G. M. Cooper. *The Cell, A Molecular Approach*. Sinauer Associates, 2nd edition, 2000.

[6] P. Gard. *Human Endocrinology*. Taylor Francis Ltd, 1998.

[7] A. Greensted and A. Tyrrell. Fault tolerance via endocrinologic based communication. In *Proceedings of ICES 2003, 5th International Conference on Evolvable Hardware*, number 2606 in LNCS, pages 24–34. Springer Verlag, March 2003.

[8] A. Greensted and A. Tyrrell. An endocrinologic-inspired hardware implementation of a multicellular system. In *Proceedings of EH 2004, 6th NASA/DoD Conference on Evolvable Hardware*, pages 245–252. IEEE Computer Society, June 2004.

[9] C. A. Janeway, P. Travers, M. Walport, and J. D. Capra. *Immunobiology, The Immune System in Health and Disease*. Garland Publishing, 4th edition, 1999.

[10] W. S. Klug and M. R. Cummings. *Concepts of Genetics*. Prentice Hall, 5th edition, 1997.

[11] P. K. Lala. *Fault Tolerant and Fault Testable Hardware Design*. Prentice Hall, 1985.

[12] H. Liu, J. Miller, and A. Tyrrell. An intrinsic robust transient fault-tolerant developmental model for digital systems. In *GECCO 2004 workshop, Genetic and Evolutionary Computation Conference*, July 2004.

[13] D. Mange, M. Sipper, A. Stauffer, and G. Tempesti. Towards robust integrated circuits: The embryonics approach. *Proceedings of the IEEE*, 88(4):516–541, April 2000.

[14] E. N. Marieb. *Human Anatomy and Physiology*. Benjamin Cummings, 5th edition, 2000.

[15] J. F. Miller. Evolving a self-repairing, self-regulating, french flag organism. In *Proceedings of GECCO 2004, Genetic and Evolutionary Computation Conference*. Springer Verlag, July 2004.

[16] C. Ortega-Sánchez and A. Tyrrell. A hardware implementation of an embryonic architecture using Virtex® FPGAs. In *Proceedings of ICES 2000, 3rd International Conference on Evolvable Hardware*, number 1801 in LNCS, pages 155–164. Springer Verlag, 2000.

[17] L. Wolpert. *The Triumph of the Embryo*. Oxford University Press, 1991.