The Knowledge Engineering Review, Vol. 21:3, 195–204. © 2006, Cambridge University Press
doi:10.1017/S0269888906000889 Printed in the United Kingdom

A health-check model for autonomic systems based on a pulse monitor

ROY STERRITT[1] and DAVE BUSTARD[2]

[1]*School of Computing and Mathematics, Jordanstown Campus, Newtownabbey, County Antrim, BT37 0QB, Northern Ireland, UK;*
e-mail: r.sterritt@ulster.ac.uk
[2]*School of Computing and Information Engineering, Coleraine Campus, Coleraine, County Londonderry, BT52 1SA, Northern Ireland, UK;*
e-mail: dw.bustard@ulster.ac.uk

Abstract

Like the autonomic responses in the human body, autonomic computing systems recognize their own *health* problems and, where possible, respond to correct them. Failing that, external help is required. The purpose of this paper is to consider how autonomic systems might be structured to facilitate health monitoring. The approach uses a 'pulse' monitor for each autonomic element, which provides a reflex reaction facility and basic information on the current state (health) of that element. The pulse mechanism extends the NASA beacon monitor concept. The different ways that pulse information might be communicated and used are examined. The discussion is illustrated with a personal computing example.

1 Introduction

Autonomic computing, first proposed by IBM in 2001 (Horn, 2001) has become established as a valuable approach to the design of robust computing systems (Mainsah, 2002; IBM, 2003). The general concepts involved are summarized in Figure 1 (Sterritt & Bustard, 2003a). An autonomic system is self-managing, meaning that it performs a range of operations to ensure its own viability. These primarily include self-protection, self-configuration, self-healing and self-optimization. Self-healing is concerned with ensuring effective recovery when a fault occurs. This means successfully identifying the fault and then, where possible, repairing it. Also, there should be minimal disruption to users, avoiding loss of data and significant delays in processing. Self-optimization means that a system is aware of its ideal performance, can measure its current performance against that ideal and has strategies for attempting improvements. A self-protecting system will defend itself from accidental or malicious external attack. This means being aware of potential threats and having ways of handling those threats. This may include self-healing actions if an attack is successful, and perhaps some self-optimization to increase protection. Finally, self-configuration is a system's ability to readjust itself automatically to changing circumstances. This may simply be in support of ongoing development or to assist in self-healing, self-optimization or self-protection. To achieve these objectives a system must be aware of its internal state (self-aware) and current external operating conditions (environment-aware). Changing circumstances, internal or external, are detected through self-monitoring and adaptations made accordingly (self-adjusting). In more detail, this means a system having knowledge of its available resources, its components, their desired performance characteristics, their current status, and the status of interconnections with other systems, along with rules and policies on how these may be adjusted. Such autonomic systems can be created through a traditional systems engineering design

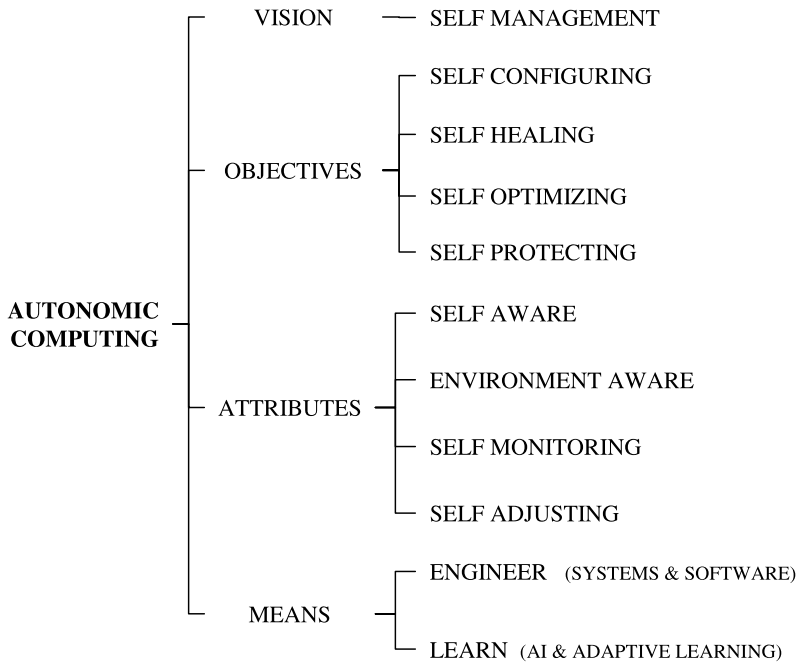

Figure 1 Autonomic computing concepts

process (Bustard *et al.*, 2005) or through inclusion of adaptive learning logic (IJCAI AI&AC Workshop, 2003).

The concept of self-managing systems has attracted significant attention from both industry and academia. In addition to the autonomic computing work at IBM, there are similar industry initiatives at HP (Adaptive Infrastructure), Sun (N1) and Microsoft (Dynamic Systems Initiative). Examples of academic involvement include work at Rutgers (active middleware) (Bhat & Parashar, 2003), CMU (self-healing systems) (Garlan *et al.*, 2003), Columbia (retrofitting legacy systems) (Kaiser *et al.*, 2002), and Imperial College (autonomic management of ubiquitous eHealth systems) (Lupu *et al.*, 2003).

Achieving the grand vision of autonomic computing is likely to involve contributions from research in many existing fields, including systems management, distributed computing, networking, operational research, software engineering, artificial intelligence, agent technology and control theory (Ganek & Corbi, 2003). Dependable and fault-tolerant computing should be especially influential, as dependability covers many system properties relevant to autonomic behaviour, such as reliability, availability, safety, security, survivability and maintainability (Avizienis *et al.*, 2000; Randell, 2000; Sterritt & Bustard, 2003a; Avizienis *et al.*, 2004).

This paper is particularly concerned with the way in which an autonomic system can provide information on its state for use by dependent autonomic elements or elements responsible for health monitoring. The next section of the paper presents a general architecture model for autonomic systems. This introduces the concept of a pulse monitor for each autonomic element through which it can report a summary of its general health. Details of the operation and use of the pulse monitor, based on the NASA beacon monitor (Wyatt *et al.*, 1998, 1999), are covered in the following section. The paper concludes with a consideration of related work which utilizes these concepts in various application areas.

2 Autonomic architecture

An autonomic system requires open standards to understand and communicate with other systems in a heterogeneous environment (Horn, 2001; IBM, 2003). Communication standards are also valuable within an autonomic system to facilitate system design and construction. For conceptual

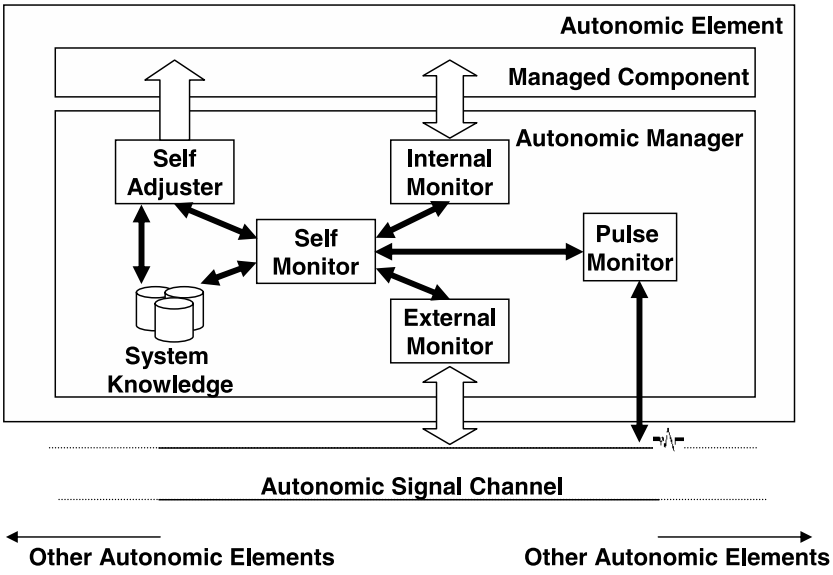

Figure 2 Abstract architecture of an autonomic element

simplicity and convenience, let us assume that an autonomic computing system is made up of a connected set of autonomic elements, each with the same basic structure and mode of operation. Figure 2 is a proposed architecture for those elements (Sterritt & Bustard, 2003b).

Here, an autonomic element is made up of a managed component and an autonomic manager. The self-monitor actively observes the state of the component (internal monitor) and its external environment (external monitor), drawing conclusions using information in the system knowledge base. If necessary, this can lead to adjustments to the managed component (self-adjuster). The external monitor observes the state of the environment through an autonomic signal channel (for instance an asynchronous event bus), which provides linkage to other autonomic managers. This linkage may be virtual (in the same physical system), peer-to-peer, client–server (Bantz *et al.*, 2003) or a Grid connection (Deen *et al.*, 2003).

A novel feature in this design is the suggested use of a pulse monitor (Sterritt, 2002, 2003b). This is used to communicate basic information about the state of each managed component to other relevant autonomic elements. In essence, its purpose is to provide a simple, general mechanism for reporting the 'health' of a managed component. This extends the biological autonomic metaphor. The design of such a pulse monitor is examined in detail in the next section.

3 Pulse monitoring

A standardized pulse monitor can provide a convenient framework for the design of fault-handling software in autonomic systems. Ideally, it should be relevant to all autonomic elements and have a simple conceptual base. One approach is to assume that the pulse monitor sends out a steady 'heartbeat'. When detected, this indicates that the autonomic element is 'alive' and when missing indicates an operational problem. It is then a small step to introduce the notion of a variable pulse rate to indicate intermediate states. NASA has used this general idea in the design of the Beacon Monitor for autonomy in space missions (Sherwood *et al.*, 1999).

3.1 NASA Beacon Monitor

NASA missions, particularly those to deep space, have had to increase onboard autonomy because of the lengthy lag-time between a craft encountering new situations and the round-trip delay in obtaining guidance from mission control (Swartwout, 1998; Wyatt *et al.*, 1998). Two of the first

missions to take this approach were DS1 (Deep Space 1) (Wyatt *et al.*, 1998) and the Mars Pathfinder (Muscettola *et al.*, 1998). The Beacon Monitor concept, first used in the DS1 mission work, automates the routine task of health monitoring, moving many responsibilities from ground to the spacecraft. The spacecraft sends a beacon signal to the ground indicating how urgent it is to track the spacecraft for telemetry. A tone is used to indicate the degree of urgency involved. Five different states are used, as follows (Sherwood *et al.*, 1999):

1. *nominal*: functions as expected; no need to establish a downlink (a communications link to download data);
2. *interesting*: interesting, non-urgent event; establish communications when convenient;
3. *important*: communications need to take place quickly or the state could deteriorate;
4. *urgent*: emergency: a critical component has failed; it is not possible to recover autonomously and intervention is needed immediately;
5. *no tone*: beacon mode is not operating.

There are a number of points to note here in relating the beacon approach to the needs of pulse monitoring in autonomic systems. The first is that there are two levels of communication involved in each case. The spacecraft uses both the beacon signal and, when necessary, direct communication, to downlink telemetry information for analysis at mission control. The same is true of autonomic systems. If an autonomic element indicates that it is operating normally, then no further action is required. If a problem is indicated, however, then the autonomic manager (through the external monitor) must be interrogated further to clarify the issue. The type of information required will be specific to the problem detected and the particular functions/services provided by the managed component.

Another similarity is that the behaviour of mission control need not be dictated by the signal from the beacon monitor. They can make direct contact with the spacecraft to obtain information whenever they wish, even if there are no apparent problems present. The same is true of communicating autonomic elements.

There are, however, some significant differences between the NASA situation and the general needs of autonomic systems. One is that interaction between the spacecraft and mission control is point-to-point, whereas many-to-many links are present in autonomic systems. Although, in principle, all autonomic elements could receive pulse signals from every other element in the system, this would quickly become overwhelming in systems of any reasonable size. The implication, therefore, is that autonomic elements need to establish explicit communication connections with each other before communication occurs. Using the biological metaphor, this can be considered equivalent to 'taking a pulse'.

Another difference is that autonomic systems have communication sequences to consider. For example, if an autonomic element detects that an element on which it relies is having problems, then it needs to report any consequential secondary problems it is experiencing. This can have a substantial knock-on ripple effect throughout an autonomic system when a serious failure occurs. It is important that this failure information be transmitted quickly to all affected elements, corresponding to the biological 'reflex reaction'. To instigate a repair, it will also be helpful to have primary and secondary problems clearly distinguished.

3.2 Pulse monitor states

Based on the requirements for pulse monitoring outlined in the preceding section, and taking account of the NASA Beacon Monitor states used, the following five-state model is proposed.

1. *Healthy*. The autonomic element is operating as required, with environment conditions as expected.
2. *Active*. The autonomic element is operating as required, but work is being done to improve its operation, triggered either by the results of an internal analysis or by recognition of environmental changes.

3. *Injured (primary and/or secondary)*. The autonomic element is experiencing difficulties affecting its performance. This may be due to internal problems (primary) and/or difficulties with its environment, such as problems with other elements on which it depends (secondary).
4. *Incapacitated (primary and/or secondary)*. The autonomic element is no longer able to function. This may be due to serious internal problems (primary) and/or extreme difficulties with its environment, such as the failure of other elements on which it depends (secondary).
5. *Dead*. The autonomic element is not communicating.

To illustrate what these states mean, consider the simple example of an autonomic element responsible for managing a file server.

1. *Healthy*. The file server is operating at the performance levels expected, saving and retrieving files successfully in all circumstances, and making security back-ups where required.
2. *Active*. The file server is making adjustments to file locations, perhaps to improve performance or in response to the addition or (controlled) removal of available storage.
3. *Injured (primary and/or secondary)*. An example of a primary problem is the file server detecting a hardware failure resulting in the loss of files. This 'injury' may be healed autonomically through retrieval of backup copies. Shortage of file space would also be treated as a primary wound. An example of a secondary problem is difficulty in retrieving files through the autonomic element managing DVD backup storage. Note that this may occur at the same time as the primary problem of file loss.
4. *Incapacitated (primary and/or secondary)*. The file server has run out of disk space. A similar problem might occur with a DVD backup management autonomic element, although this alone would only 'injure' the file server.
5. *Dead*. The file server cannot be contacted because, for example, the server software has crashed, the server host machine has failed, or the server has been disconnected from the network unexpectedly. The next two subsections consider how the pulse monitor might be implemented and illustrate its use in a personal computing system, expanding on the file server example.

3.3 Pulse monitor implementation

The NASA Beacon Monitor uses a variable tone to indicate its different states. An autonomic pulse monitor could use a similar mechanism, indicating different states through different frequencies of communication. This has the advantage of being consistent with the biological metaphor, in which 'no signal' means 'death' and higher pulse rates indicate greater degrees of stress. In computing terms, however, each signal can easily include extra information to avoid the receiving element having to deduce a state change. Also, once the receiving element has acknowledged receipt of a problem, there seems little point in repeating the message frequently, consuming processing power that might be needed to provide assistance. It is still important, however, to provide the basic carrier heartbeat to be able to recognise quickly when an autonomic element has 'died'. This must be communicated rapidly through the system to avoid further problems and perhaps contact human operators who may be needed to resolve the problem.

3.4 Personal computing example

To illustrate the way that the 'health' of an autonomic system might be monitored, consider the following personal computing example. A local network is set up as shown in Figure 3. This comprises autonomic elements for:

i. each connected user;
ii. a shared file server;
iii. a shared printer server;
iv. a DVD backup server;
v. an operator with overall responsibility for the network.

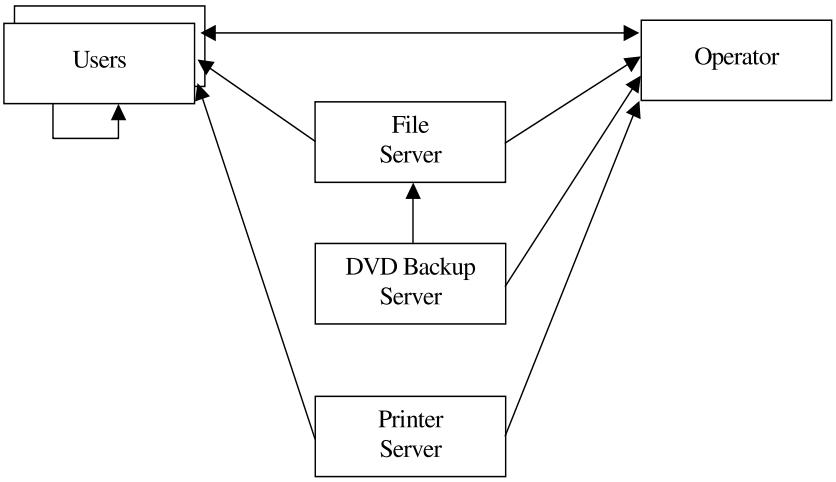

Figure 3 Pulse monitor connections in a personal computing system

The lines in the diagram indicate monitoring relationships between autonomic elements. These indicate, for example, that each user monitors the printer server, the operator monitors the users, and the users monitor the operator.

In general, there are a number of ways in which the 'health' of autonomic elements might be monitored. One is to have dedicated system elements for that purpose. Another, more robust approach, is to share this responsibility among the autonomic elements so that, in effect, they monitor each other. It is convenient, for example, for an autonomic element to monitor the health of any other system element on which it depends. To improve robustness it is also desirable to have every autonomic element monitored at least twice, as indicated in the example. The users in this case might be connected in a two-way ring, with each monitoring their left and right neighbour.

Most autonomic research, especially that in industry, has so far focused on server management (Bantz & Frank, 2003), since requirements for reliability have increased, while servers have become more complex and hence more difficult to maintain. The contribution of autonomic computing to personal computing is different being much less about achieving optimum performance or exploiting redundancy and more about simplifying use of the equipment and the associated services involved (Bantz & Frank, 2003; Sterritt & Bantz, 2004). This is an important area because, potentially, it can affect every computer user. Issues to be considered include, for example:

- how to design file servers to help users become better organized, and hence make better use of the storage space available;
- how to deal effectively with laptop connections to a network to facilitate users making the connection while protecting the network from viruses and other threats;
- how to involve users in autonomic activity, using their technical knowledge to recover from operational problems (e.g. a hardware problem while the operator is unavailable) or provide assistance for other users;
- how to design autonomic elements to respond effectively to system problems that occur. Certainly, it is important for elements to be informed quickly (reflex reaction) but what action should then follow? As in biological systems, a short-term reflex action coupled with a longer-term healing process seems the necessary strategy to adopt (Bapty *et al.*, 2003). Human reflex reactions enable a rapid response to pain, such as when a hot object is touched. In computing terms, it is likely that each autonomic system will have to reconfigure itself in response to problems encountered, while maintaining its overall operation as far as possible. This may result in the system operating with a reduced set of resources and services (Bapty *et al.*, 2003).

4 Related work

The section briefly discusses two application areas where the pulse monitor may be considered beneficial.

4.1 Grid Heartbeat Monitor

The Globus Open Grid Services Architecture (OGSA) has a health-check facility known as the Globus Heartbeat Monitor, which is designed to provide a simple, highly reliable mechanism for monitoring processes (Stelling *et al.*, 1998). The Globus OGSA specifies three aspects of the heartbeat monitor: a client library, a local monitor, and a data collector. Essentially, a local monitor runs on each host checking and reporting on the status of its local processes and the overall system, generating 'I-am-alive' messages (heartbeats) to confirm working health.

The client library enables processes to register when they are activated within the local monitor and sign out when they terminate. Each local monitor periodically executes a review cycle in which it checks the status of the client processes it is monitoring, updates the local status information of those processes, and sends a report on each process to one or more external agents (data collectors) specified at registration. There can be any number of data collectors defined, but typically there is at least one for tracking all of the monitored processes associated with the computing environment, and one for each distributed application. A data collector infers that monitored components have failed or are unavailable, based on reports it is expecting from local monitors but has not received (missing heartbeats) (Stelling *et al.*, 1998).

This mechanism could usefully be extended with the pulse concept to provide more information on operational health. For instance, a local monitor recognizing that a process is no longer executing optimally might send this health indicator pulse to the data collector. The data collector would then have the benefit of additional information in assessing the heath of a service and be able to detect deterioration prior to complete failure, taking remedial action as appropriate to reduce the risk identified (Sterritt 2003b).

4.2 Fault management in telecommunication systems

These approaches have been investigated in relation to more effective fault management in telecommunication systems, particularly in exploring the relationship between a rapid response to faults that occur (reflex reaction) and the slower analysis and healing process that follows. Figure 4 shows a typical fault management architecture for a large telecommunications network (Sterritt *et al.*, 2004). Faults occur in network elements at the bottom of the reporting tree in the figure. Many different network technologies are involved (e.g. SDH (SONET), PDH, ATM, IP) and each has a specific fault manager. The network technologies are interdependent to some extent, however, so it is only at the next layer up, the cross technology network fault manager, that the faults reported can be correlated to produce a coherent understanding of the root cause of the problem.

Once the root cause has been determined, either automatically or through operator assistance, the fault is assigned a trouble ticket and handed over to field force management for attention. Because of the amount of processing involved in the various layers shown, it can take up to 15 minutes for details of a fault to reach the service layer, by which time customers may have already reported the problem.

A study took in 2003 on the use of autonomicity in telecommunications at BT (Sterritt, 2003a; Sterritt *et al.*, 2004, 2005) looked specifically at the problem of slow fault reporting. It was found that such delays could, in principle, be reduced substantially by having pulse monitors in each (autonomic) fault management element that can supply reflex health signals about fault traffic ahead of details on the exact source of the problem emerging through the hierarchy. This may be achieved through extending the existing heartbeat monitors between the managers to carry a network health indicator (the pulse), as depicted in Figure 4.

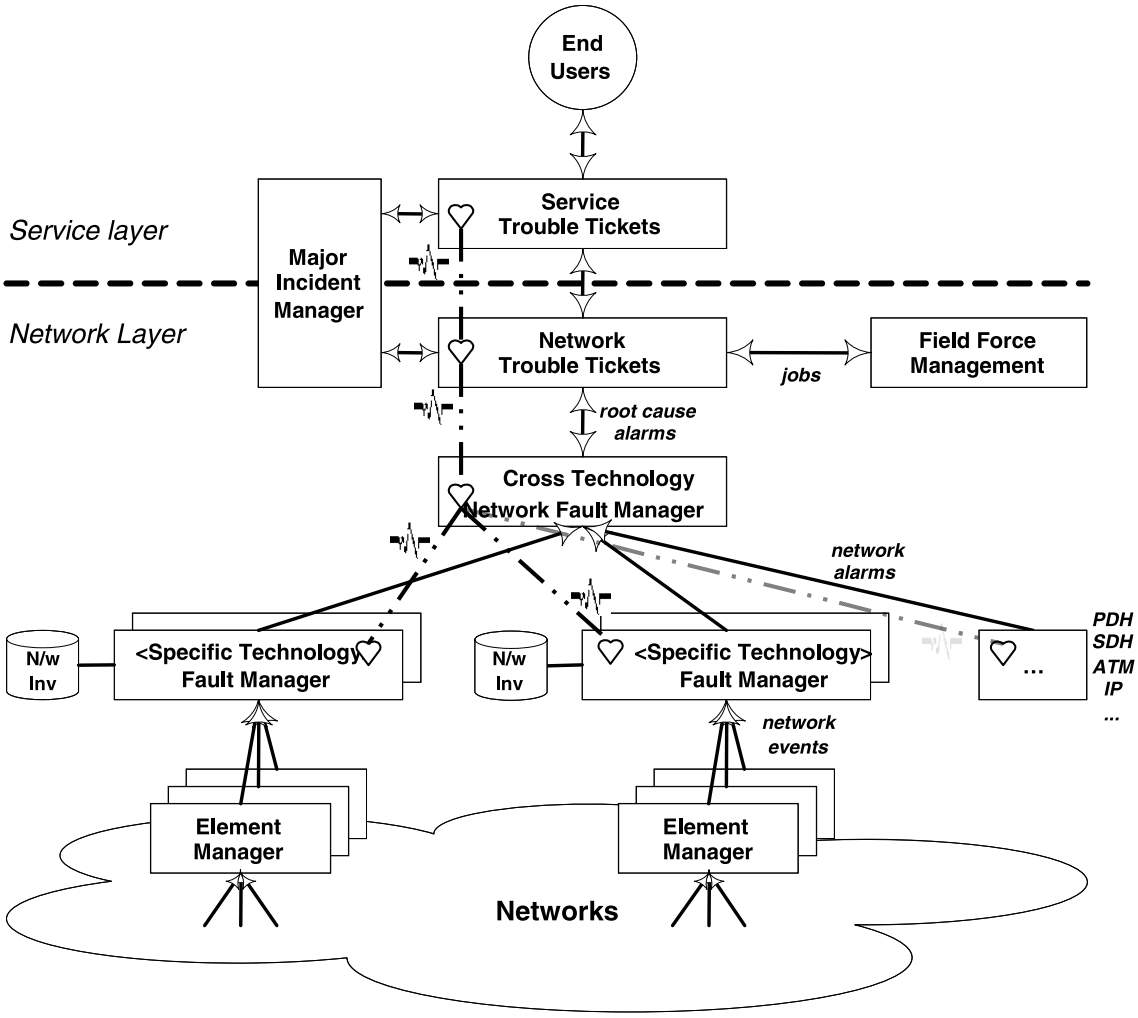

Figure 4 Large telecommunications fault management architecture

5 Conclusions

Potentially, the autonomic computing concept could have an impact on system design similar to that of the object-oriented paradigm. It promises to improve the robustness of future systems, as well as reducing complexity for users, by taking on more responsibility for identifying and handling system problems that arise. The purpose of this paper was to consider how system 'health' might be monitored. This was based on an assumption that autonomic systems were made up of a set of autonomic elements with the same basic architecture. Included in that architecture was a pulse monitor, to provide a reflex reaction and essential information on the current state of health of each autonomic element. This is an extension of the more commonly used 'heartbeat' mechanism, typically found in embedded systems and fault-tolerant computing applications.

Each pulse monitor can report one of five different states for its associated element: healthy, active, injured, incapacitated and dead. These are a variation of the states developed for the NASA Beacon Monitor, adapted to cover general computing and communications health summarization. The use of the states was illustrated with a file server autonomic element, operating in a personal computing network. The discussion identified a number of areas for future research, particularly in personal computing. Overall, however, it seems possible that such a mechanism could be of value in the construction of almost any system.

Acknowledgements

This work was undertaken through the Centre for Software Process Technologies, which is supported by the EU Programme for Peace and Reconciliation in Northern Ireland and the Border Region of Ireland (PEACE II).

References

Avizienis, A., Laprie, J.-C. & Randell, B. 2000 Fundamental concepts of dependability. UCLA CSD Report no. 010028.

Avizienis, A., Laprie, J.-C., Randell, B. & Landwehr, C. 2004 Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing* **1**(1), 11–33.

Bantz, D. F., Bisdikian, C., Challener, D., Karidis, J. P., Mastrianni, S., Mohindra, A., Shea, D. G. & Vanover, M. 2003 Autonomic personal computing. *IBM Systems Journal* **42**(1), 165–176.

Bantz, D. F. & Frank, D. 2003 Challenges in autonomic personal computing, with some new results in automatic configuration management. In *Proceedings of IEEE Workshop on Autonomic Computing Principles and Architectures (AUCOPA 2003) at INDIN 2003, Banff, Alberta, Canada, 22–23 August*.

Bapty, T., Neema, S., Nordstorm, S., Shetty, S., Vashishtha, D., Overdorf, J. & Sheldon, P. 2003 Modeling and generation tools for large-scale, real-time embedded systems. In *Proceedings of the 10th IEEE International Conference and Workshop on the Engineering of Computer-Based Systems (ECBS), Huntsville, AL, 7–10 April*, pp. 11–16.

Bhat, V. & Parashar, M., 2003 Discover middleware substrate for integrating services on the Grid. In *Proceedings of the 10th International Conference on High Performance Computing (HiPC '03), Hyderabad, India, December*.

Bustard, D. W., Sterritt, R., Taleb-Bendiab, A., Laws, A., Randles, M. & Keenan, F. 2005 Towards a systemic approach to autonomic systems engineering. In *Proceedings of the IEEE Workshop on the Engineering of Autonomic Systems (EASe 2005) at 12th Annual IEEE International Conference and Workshop on the Engineering of Computer Based Systems (ECBS 2005), Greenbelt, MD, 3–8 April*, pp. 465–472.

Deen, G., Lehman, T. & Kaufman, J. 2003 The Almaden optimalGrid project. In *Proceedings of the Autonomic Computing Workshop, 5th International Workshop on Active Middleware Services (AMS 2003), Seattle, WA*, pp. 14–21.

Ganek, A. G. & Corbi, T. A. 2003 The dawning of the autonomic computing era. *IBM Systems Journal* **42**(1), 5–18.

Garlan, D., Shang-Wen Cheng & Bradley Schmerl 2003 Increasing system dependability through architecture-based self-repair. In *Architecting Dependable Systems*, de Lemos, R. Gacek, C. Romanovsky BA. (eds.). Springer.

Horn, P. 2001 Autonomic computing: IBM perspective on the state of information technology, IBM T. J. Watson Labs, NY, 15 October 2001. *Presented at AGENDA 2001, Scotsdale, AR.*

IBM 2003 An architectural blueprint for autonomic computing. White Paper, release: April'03.

Kaiser, G., Gross, P., Kc, G., Parekh, J. & Giuseppe Valetto, 2002 An approach to autonomizing legacy systems. In *Proceedings of the Workshop on Self-Healing, Adaptive and Self-MANaged Systems, June*.

Lupu, E., Sloman, M., Dulay, N. & Sventek, J. 2003 AMUSE: Autonomic management of ubiquitous systems for e-health. EPSRC project, available at `http://www.doc.ic.ac.uk/~ecl1/projects/AMUSE/`.

Mainsah, E., 2002 Autonomic computing: The next era of computing. *IEE Electronics Communication Engineering Journal* **14**(1), 2–3.

Muscettola, N., Nayak, P. P., Pell, B. & Williams, B. 1998 Remote agent: to boldly go where no AI system has gone before. *Artificial Intelligence* **103**(1–2), 5–48.

IJCAI Workshop 2003 AI and autonomic computing: developing a research agenda for self-managing computer systems. In *Proceedings, Acapulco, Mexico, 10 August*. Available at `http://www.research.ibm.com/ACworkshop`.

Randell, B. 2000 Turing memorial lecture facing up to faults. *Computer Journal* **43**(2), 95–106.

Sherwood, R., Wyatt, J., Hotz, H., Schlutsmeyer, A. & Sue, M. 1999 Lessons learned during implementation and early operations of the DS1 Beacon monitor experiment. In *Proceedings of the 3rd International Symposium on Reducing the Cost of Ground Systems and Spacecraft Operations, Tainan, Taiwan*.

Stelling, P. F., Foster, I. & Kesselman, C. 1998 Fault detection services for wide area distributed computations. *Retreat*.

Sterritt, R. 2002 Towards autonomic computing: Effective event management. In *Proceedings of the 27th Annual IEEE/NASA Software Engineering Workshop*. Los Alamitos, CA: IEEE Computer Society Press, pp. 40–47.

Sterritt, R. 2003a xACT: Autonomic computing and telecommunications. *British Telecom Research Fellowship*.

Sterritt R. 2003b Pulse monitoring: extending the health-check for the autonomic GRID. In *Proceedings of the IEEE Workshop on Autonomic Computing Principles and Architectures (AUCOPA' 2003) at IEEE International Conference Industrial Informatics (INDIN 2003), Banff, Alberta, Canada, 22–23 August*.

Sterritt, R. & Bantz, D. F. 2004 PAC-MEN: personal autonomic computing monitoring environments. In *Proceedings of IEEE DEXA 2004 Workshops — 2nd International Workshop on Self-Adaptive and Autonomic Computing Systems (SAACS 04), Zaragoza, Spain, 30 August–3 September*. Los Alamitos, CA: IEEE Computer Society Press, pp. 737–741.

Sterritt, R. & Bustard, D. W. 2003a Autonomic computing — a means of achieving dependability? In *Proceedings of IEEE International Conference on the Engineering of Computer Based Systems (ECBS'03), Huntsville, Alabama, 7–11 April*, pp. 247–251.

Sterritt, R. & Bustard, D. W. 2003b Towards an autonomic computing environment. In *Proceedings of the 1st International Workshop on Autonomic Computing Systems at 14th International Conference on Database and Expert Systems Applications (DEXA'2003), Prague, Czech Republic, 1–5 September*.

Sterritt, R., Gunning, D., Meban, A. & Henning, P. 2004 Exploring autonomic options in an unified fault management architecture through reflex reactions via pulse monitoring. In *Proceedings of IEEE Workshop on the Engineering of Autonomic Systems (EASe 2004) at the 11th Annual IEEE International Conference and Workshop on the Engineering of Computer Based Systems (ECBS 2004), Brno, Czech Republic, 24–27 May*, pp. 449–455.

Sterritt, R., Bustard, D. W., Gunning, D. & Henning, P. 2005 Autonomic communications and the reflex unified fault management architecture. *Advanced Engineering Informatics* **19**, 189–198.

Swartwout, M. A. 1998 Engineering data summaries for space missions. *SSDL*.

Wyatt, J., Hotz, H., Sherwood, R., Szijjaro, J. & Sue, M. 1998 Beacon monitor operations on the deep space one mission. In *Proceedings of the 5th International Symposium Robotics and Automation in Space, Tokyo, Japan*.

Wyatt, J., Sherwood, R., Sue, M. & Szijjarto, J. 1999 Flight validation of on-demand operations: the deep space one beacon monitor Operations Experiment. In *Proceedings of the 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS '99)*, ESTEC, *Noordwijk, The Netherlands, 1–3 June*.