

Towards Self-Managing Real-Time Systems

Roy Sterritt
University of Ulster
Jordanstown,
Northern Ireland
Email: r.sterritt@ulster.ac.uk

Mike Hinchey
Lero—the Irish Software Engineering Research
Centre, University of Limerick, Ireland
mike.hinchey@lero.ie.

Abstract—Upon first consideration, the Self-Managing Systems (and specifically Autonomic Systems) paradigm may seem to add too much additional functionality and overhead for systems that require strict real-time behavior. From our experience, not only does the Real-Time Systems area benefit from autonomicity, but the autonomic systems initiative requires the expertise of the RTS community in order to achieve its overarching vision. We briefly highlight some of the design considerations we have taken in incorporating self-management into Real-Time Systems.

I. INTRODUCTION

“AUTONOMIC Computing” is the term coined, initially by DARPA under the Autonomic Information Assurance program [13], and subsequently by IBM in their call to industry in 2001 [2]. The approach takes inspiration from the human and mammalian nervous systems in attempting to develop systems that are self-managing and ultimately self-governing [1].

We have argued in a previous article [3] that computer-based systems should be autonomic. We stand by that belief, in particular as software becomes more complex with greater expectations in terms of functionality, performance and real-time behavior. In our previous article [3], we were addressing a particular audience, viz. members of the IEEE Technical Committee on Engineering Computer-Based Systems (TC-ECBS). To that community, “computer-based systems” typically means embedded systems involving both hardware and software.

Such systems, like real-time systems, often have strict constraints in terms of performance and timing requirements. The addition of concepts from Autonomic Computing (AC) might at a first glance seem to add complexity rather than reduce it, something that real-time systems can generally ill-afford. We believe, however, that Autonomic Computing has much to offer in terms of reducing complexity, and our work over the last number of years has looked at how we might exploit AC and other biologically-inspired techniques in classes of systems that would otherwise be infeasible.

We believe that Autonomic Computing is not achievable without real-time systems (in particular at lower, implementation levels). Simultaneously, we believe that real-time systems can benefit much from AC, without significant overhead, and that future advancements necessitate that we move towards self-managing real-time systems.

II. SELF-MANAGING SYSTEMS

The vision of achieving self-management in our computing and communications systems has come to light under many similar initiatives—Autonomic Computing, Autonomic Communications, Organic Computing, Adaptive Infrastructure, N1, Dynamic System Initiative, Adaptive Network Care, Proactive Computing, Self-Organizing Systems, Self-*, Selfware, Biologically-Inspired Computing [5].

Although these different initiatives may have different academic specialties at their base, or were initiated by different industrial organizations, often they are inspired by self-management as exhibited in Biology and Nature. The IEEE has formalized the community as a Technical Committee on “Autonomous and Autonomic Systems” to underscore the strategic aims of developing next generation self-organizing and self-managing computing and communications infrastructures and systems.

A. Autonomic Computing

The Autonomic Computing (AC) initiative focuses on managing complexity with self-managing systems, taking inspiration from the human autonomic nervous system (ANS) [1]-[5].

The ANS is that part of the nervous system that controls the vegetative functions of the body, such as circulation of the blood, intestinal activity, and secretion and production of chemical “messengers” (hormones) that circulate in the blood. The sympathetic nervous system (SyNS) supports “fight or flight”, providing various protection mechanisms to ensure the safety and well-being of the body. The parasympathetic nervous system (PaNS) supports “rest and digest”, ensuring that the body performs necessary functions for long-term health.

The general properties of an autonomic (self-managing) system can be summarized by four objectives: being self-configuring, self-healing, self-optimizing and self-protecting, and four attributes: self-awareness, self-situated, self-monitoring and self-adjusting [7]. Essentially, the objectives represent broad system requirements, while the attributes identify basic implementation mechanisms.

Self-configuring represents a system’s ability to re-adjust itself automatically; this may simply be in support of changing circumstances, or to assist in self-healing, self-optimization or self-protection.

Self-healing, in reactive mode, is a mechanism concerned with ensuring effective recovery when a fault occurs, identifying the fault, and then, where possible, repairing it. In proactive mode, it monitors vital signs in an attempt to predict and avoid “health” problems (reaching undesirable situations).

Self-optimization means that a system is aware of its ideal performance, can measure its current performance against that ideal, and has defined policies for attempting improvements. It may also react to policy changes within the system as indicated by the users. A self-protecting system will defend itself from accidental or malicious external attack. This necessitates awareness of potential threats and a means of handling those threats.

In achieving such self-managing objectives, a system must be aware of its internal state (self-aware) and current external operating conditions (self-situated). Changing circumstances are detected through self-monitoring and adaptations are made accordingly (self-adjusting).

As such, a system must have knowledge of its available resources, its components, their desired performance characteristics, their current status, and the status of inter-connections with other systems, along with rules and policies for how these may be adjusted. Such ability to operate in a heterogeneous environment will require the use of open standards to enable global understanding and communication with other systems.

These mechanisms are not independent entities. For instance, if an attack is successful, this will necessitate self-healing actions, and a mix of self-configuration and self-optimization, in the first instance to ensure dependability and continued operation of the system, and later to increase self-protection against similar future attacks. Finally, these self-mechanisms should ensure that there is minimal disruption to users, avoiding significant delays in processing.

The Autonomic Systems architecture essentially consists of cooperating autonomic elements made up of the component that is required to be managed, and the autonomic manager [11, 12]. It is assumed that an autonomic manager (AM) is responsible for a managed component (MC) within a self-contained autonomic element (AE). This autonomic manager may be designed as part of the component or provided externally to the component, as an agent for instance. To achieve a self-managing system these AEs will cooperate with remote autonomic managers through virtual, peer-to-peer, client-server or grid configurations.

B. Aims of AC

If we consider the eight overriding objectives set out by IBM in 2001 in their call to the industry [2], some of these will seem, to the real-time systems (RTS) researcher or practitioner, characteristics that are common to real-time systems or that at least very desirable:

1. To be autonomic, a computing system needs to “know itself”—and comprise components that also possess a system identity.
2. An Autonomic Computing System must configure and reconfigure itself under varying and unpredictable conditions.

3. An Autonomic Computing system never settles for the status quo—it always looks for ways to optimize its workings.
4. An Autonomic Computing system must perform something akin to healing—it must be able to recover from routine and extraordinary events that might cause some of its parts to malfunction.
5. A virtual world is no less dangerous than the physical one, so an Autonomic Computing system must be an expert in self-protection.
6. An Autonomic Computing system knows its environment and the context surrounding its activity, and acts accordingly.
7. An Autonomic Computing system cannot exist in a hermetic environment.
8. *Perhaps* most critical for the user, an Autonomic Computing system will anticipate the optimized resources needed while keeping its complexity hidden.

This list of eight requirements, set out in [2] and [4], point to the fact that to qualify as an Autonomic System (AS), a system has to have a degree of self-awareness and familiarization with the components that compose it. It must know their capabilities, if it is to be able to adapt to its environment successfully, analogous to the way that dynamic scheduling requires a knowledge of tasks, their periods and deadlines.

An AS must be able to adapt to unpredictable conditions and adjust itself as necessary. This is something that is done, almost routinely, in an RTS. Real-Time Systems are constantly evolving. Their role is to adapt to changing environmental circumstances or new inputs. The role of scheduling is to optimize the use of resources and to make efficient use of available processing power. An effective RTS is performing this role.

A key focus of the list of essentials given above, is a move towards management and reduction of inherent complexity. Complexity is a major issue for any large-scale system. RTS have the added difficulty of having to meet strict timing-constraints and inputs from many sources, all of which need to be met.

But surely adding self-managing abilities to an already complex system, such as exhibited by an RTS, is only adding to inherent complexity? The overhead of self-management essentially adds an overhead to any system. Many RTSs are already pushing the limits of their schedules and adding any additional overhead is impossible.

It is our contention, however, that sometimes adding this overhead is justified. There are classes of systems for which self-management, and even self-government, is essential, for which timing constraints could not be met otherwise and for which the overhead of AC is more than repaid by a corresponding reduction in complexity.

C. Implementing AC

It will be obvious to any practitioner that to hand over control to the system that was previously held by a human, in principle will add additional functionality, complexity and processing overhead to the system. Also experience from AI, Expert Systems, and Machine Learning problems will

have shown many how slow decision-making systems can be. As such, it is vital to develop key principles and engineering techniques for how self-management systems should be created.

If we reconsider the meaning of the actual terms “Autonomic” and “Autonomous” [6]:

au-to-nom-ic (àwtə nómmik)

adj.

1. Physiology.
 - a. Of, relating to, or controlled by the autonomic nervous system.
 - b. Occurring involuntarily; automatic: an autonomic reflex.
2. Resulting from internal stimuli; spontaneous.

au-ton-o-mic-i-ty. (àwtə nóm i síttee)

n.

1. The state of being autonomic.

au-ton-o-mous (aw tónnəməs)

adj.

1. Not controlled by others or by outside forces; independent: an autonomous judiciary; an autonomous division of a corporate conglomerate.
2. Independent in mind or judgment; self-directed.
3.
 - a. Independent of the laws of another state or government; self-governing.
 - b. Of or relating to a self-governing entity: an autonomous legislature.
 - c. Self-governing with respect to local or internal affairs: an autonomous region of a country.
4. Autonomic.

[From Greek *autonomos*: *auto-*, *auto-* + *nomos*, *law*]

A difference stands out: “Autonomic” is very much concerned with spontaneous, reflex reactions while “Autonomous” is a slower, high level conscious decision-making process.

The basic principles of Autonomic and Autonomous systems can be incorporated into the design of a system to ensure that the correct response rate is achieved where it is needed. This has resulted in us considering a simple three tiered abstract architecture in our designs of self-managing systems:

- Autonomous Layer
- Selfware Layer
- Autonomic Layer

The *Autonomic Layer* is the bottom tier, closest to the hardware, and operates with immediate reaction to situations to ensure that system operations are maintained.

The *Selfware Layer* incorporates day-to-day operations of self-managing activity as and when needed, and as and when the system has the processing capacity available.

The *Autonomous Layer* is the top tier where high-level strategic objectives of the system are directed and satisfied over time. This often includes reflection.

As such, a key element that we have included in our work in designing AS is that of the *Autonomic Reflex*, borrowed

from embedded and real-time systems and extended to include active system health telemetry.

D. Reflex Reactions

Essentially, the aim of Autonomic Computing is to create robust, dependable self-managing systems [8] in an attempt to deal with complexity.

At the heart of the architecture of any autonomic system are sensors and effectors. A control loop is created by monitoring behavior through sensors, comparing this with expectations (knowledge, as in historical and current data, rules and beliefs), planning what action is necessary (if any), and then executing that action through effectors. The closed loop of feedback control provides the basic backbone structure for each system component. There are two conceptual control loops in an Autonomic Element—one for self-awareness and another for self-situation (environmental awareness and context-awareness).

IBM represents this self-monitor/self-adjuster control loop as the monitor, analyze, plan and execute (MAPE) control loop. The monitor-and-analyze parts of the structure process information from the sensors to provide both self-awareness and an awareness of the external environment. The plan-and-execute parts decide on the necessary self-management behavior that will be executed through the effectors. The MAPE components use the correlations, rules, beliefs, expectations, histories, and other information known to the autonomic element, or available to it through the knowledge repository within the AM.

The autonomic environment requires that autonomic elements and, in particular, autonomic managers communicate with one another concerning self-* activities, in order to ensure the robustness of the environment [9, 10]. It is our belief that the autonomic manager communications (AM ⇔ AM) must also include a reflex signal.

To facilitate this, fault-tolerant mechanisms such as a heart-beat monitor (‘I am alive’ signals) and pulse monitor (urgency/reflex signals) may be included within the autonomic element [9, 10]. See Figure 1 for an illustration of our autonomic environment.

The notion behind the pulse monitor (PBM) is to provide an early warning of an undesirable condition so that preparations can be made to handle the processing load of diagnosis and planning a response, including diversion of load. Together with other forms of communications it creates dynamics of autonomic responses [11] — the introduction of multiple loops of control, some slow and precise, others fast and possibly imprecise, fitting with the biological metaphors of reflex and healing [9].

This reflex component may be used to safeguard the autonomic element by communicating its health to another AE. The component may also be utilized to communicate environmental health information. For instance, in the situation where each PC in a LAN is equipped with an autonomic manager, rather than each of the individual PCs monitoring the same environment, a few PCs (likely the pillow queen—the least busy machines) may take on this role and alert the others via a change in pulse to indicate changing circumstances.

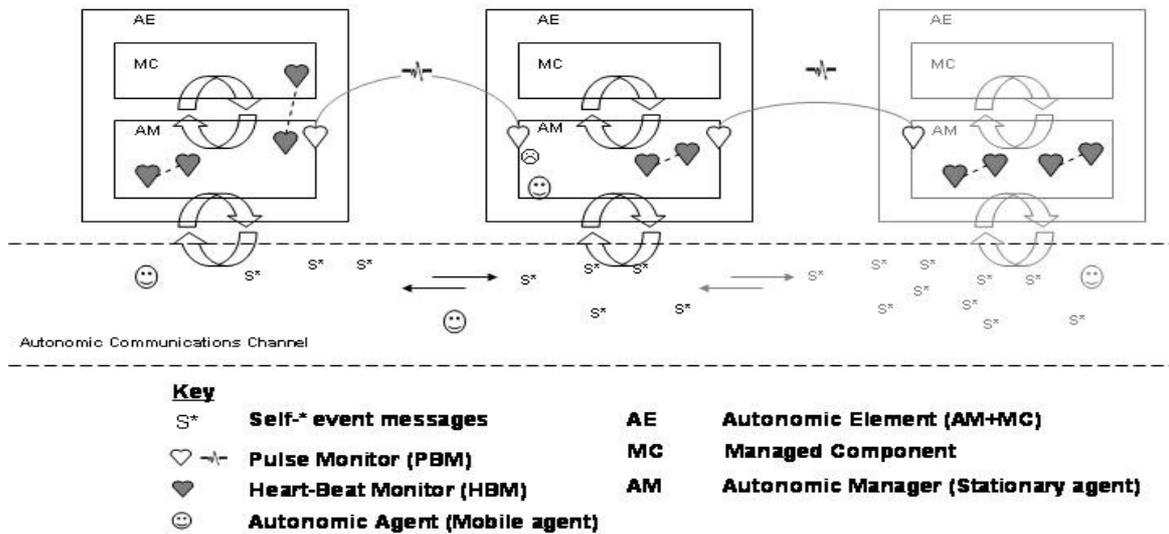


Figure 1: An Autonomic Environment

An important aspect concerning the reflex reaction and the pulse monitor is the minimization of data sent—essentially only a “signal” is transmitted. Strictly speaking, this is not mandatory; more information may be sent, yet the additional information must not compromise the reflex reaction and the required real-time response of the system. For instance, in the absence of bandwidth concerns, information that can be acted upon quickly and not incur processing delays could be sent. The important aspect is that the information must be in a form that can be acted upon immediately and not involve processing delays (such as is the case of event correlation).

Just as the beat of the heart has a double beat (lu b-dub) the autonomic element’s pulse monitor may have a double beat encoded as described above, a self health/urgency measure and an environment health/urgency measure. These match directly with the two control loops within the AE, and the self-awareness and environment awareness properties.

Another extension we have been working on is the adaptive pulse monitor; where the rate of the pulse adapts not only when alerting other self-managing components about concerns (increase in pulse rate) but to take into consideration bandwidth concerns. In particular, under fault conditions there is often a cascade and flood of event and alarm messages on the network, and we must actively reduce alerting so that achieving autonomicity is not making the situation worse.

III. A SELF-MANAGING RTS EXAMPLE

IV. NASA Missions

NASA missions require the use of complex hardware and software systems, and embedded systems, often with hard real-time requirements [3]. Most missions involve significant degrees of autonomous behavior, often over significant periods of time. There are missions which are intended only to survive for a short period, and others which will continue for decades, with periodic updates to both hardware and software. Some of these updates are pre-planned; others, such as with the Hubble Space Telescope, were not planned but now will be undertaken (with updates performed either by astronauts or via a robotic arm).

While missions typically have human monitors, many missions involve very little human intervention, and then often only in extreme circumstances. It has been argued that NASA systems should be autonomic [9, 14], and that all autonomous systems should be autonomic by necessity. Indeed, the trend is in that direction in forthcoming NASA missions.

We take as our example, a NASA concept mission, ANTS, which has been identified [15] as a prime example of an autonomic system.

1) ANTS

ANTS is a concept mission that involves the use of intelligent swarms of spacecraft. From a suitable point in space (called a Lagrangian), 1000 small spacecraft will be launched towards the asteroid belt.

As many as 60% to 70% of these will be destroyed immediately on reaching the asteroid belt. Those that survive will coordinate into groups, under the control of a leader, which will make decisions for future investigations of particular asteroids based on the results returned to it by individual craft which are equipped with various types of instruments

a) Self-configuring

ANTS will continue to prospect thousands of asteroids per year with large but limited resources. It is estimated that there will be approximately one month of optimal science operations at each asteroid prospected. A full suite of scientific instruments will be deployed at each asteroid. ANTS resources will be configured and re-configured to support concurrent operations at hundreds of asteroids over a period of time.

The overall ANTS mission architecture calls for specialized spacecraft that support division of labor (rulers, messengers) and optimal operations by specialists (workers). A major feature of the architecture is support for cooperation among the spacecraft to achieve mission goals. The architecture supports swarm-level mission-directed behaviors, sub-swarm levels for regional coverage and resource-sharing, team/worker groups for coordinated science operations and individual autonomous behaviors. These organizational levels are not static but evolve and self-configure as the need arises. As asteroids of interest are identified, appropriate teams of spacecraft are configured to realize optimal science operations at the asteroids. When the science operations are completed, the team disperses for possible reconfiguration at another asteroid site. This process of configuring and reconfiguring continues throughout the life of the ANTS mission.

Reconfiguring may also be required as the result of a failure, such as the loss of, or damage to, a worker due to collision with an asteroid (in which case the role may be assumed by another worker, which will be allocated the task and resources of the original).

b) Self-healing

ANTS is self-healing not only in that it can recover from mistakes, but self-healing in that it can recover from failure, including damage from outside forces. In the case of ANTS, these are non-malicious sources: collision with an asteroid, or another spacecraft, etc.

ANTS mission self-healing scenarios span the range from negligible to severe. A negligible example would be where an instrument is damaged due to a collision or is malfunctioning. In such a scenario, the self-healing behavior would be the simple action of deleting the instrument from the list of functioning instruments. A severe example would arise when the team loses so many workers it can no longer conduct science operations. In this case, the self-healing behavior would include advising the mission control center and requesting the launch of replacement spacecraft, which would be incorporated into the team, which in turn would initiate necessary self-configuration and self-optimization.

Individual ANTS spacecraft will have self-healing capabilities also. For example, an individual may have the capability of detecting corrupted code (software), causing it to re-

quest a copy of the affected software from another individual in the team, enabling the corrupted spacecraft to restore itself to a known operational state.

c) Self-optimizing

Optimization of ANTS is performed at the individual level as well as at the system level.

Optimization at the ruler level is primarily through learning. Over time, rulers will collect data on different types of asteroids and will be able to determine which asteroids are of interest, and which are too difficult to orbit or collect data from. This provides optimization in that the system will not waste time on asteroids that are not of interest, or endanger spacecraft examining asteroids that are too dangerous to orbit.

Optimization for messengers is achieved through positioning, in that messengers may constantly adjust their positioning in order to provide reliable communications between rulers and workers, as well as with mission control back on Earth.

Optimization at the worker level is again achieved through learning, as workers may automatically skip over asteroids that it can determine will not be of interest.

d) Self-protecting

The significant causes of failure in ANTS will be collisions (with both asteroids and other spacecraft), and solar storms.

Collision avoidance through maneuvering is a major challenge for the ANTS mission, and is still under development. Clearly there will be opportunity for individual ANTS spacecraft to coordinate with other spacecraft to adjust their orbits and trajectories as appropriate. Avoiding asteroids is a more significant problem due to the highly dynamic trajectories of the objects in the asteroid belt. Significant planning will be required to avoid putting spacecraft in the path of asteroids and other spacecraft.

In addition, charged particles from solar storms could subject spacecraft to degradation of sensors and electronic components. The increased solar wind from solar storms could also affect the orbits and trajectories of the ANTS individuals and thereby could jeopardize the mission. One possible self-protection mechanism would involve a capability of the ruler to receive a warning message from the mission control center on Earth. An alternative mechanism would be to provide the ruler with a solar storm sensing capability through on-board, direct observation of the solar disk. When the ruler recognizes that a solar storm threat exists, the ruler would invoke its goal to protect the mission from harm from the effects of the solar storm, and issue instructions for each spacecraft to "fold" the solar sail (panel) is uses to charge its power sources.

e) Self-aware

Clearly, the above properties require the ANTS mission to be both aware of its environment and self-aware.

The system must be aware of the positions and trajectories of other spacecraft in the mission, of positions of asteroids

and their trajectories, as well as of the status of instruments and solar sails.

V. Real Time Issues

The swarm-based concepts of ANTS (or its submission, PAM—Prospecting Asteroid Mission—as described above) enable exploration missions that never before would be possible. Such concept missions are clearly real-time systems. ANTS must be survivable in a harsh environment (space) over multiple years. The mission must be able to protect itself and to recover from collisions, threats from solar storms, and other problematic issues.

This must all be considered in the context of significant transmission delays. Round-trip delays between Earth and the asteroid belt exceed 40 minutes. The result is that exceptional events cannot be dealt with from Earth. Even anticipated events cannot be dealt with from Earth, as catastrophic damage could have occurred before ground control had even received notification.

The result is that the system must be self-managing. In order for its real-time behavior to be realized, the mission must exhibit the properties of an Autonomic System, which (as we pointed out in Section II.A) are desirable properties of a RTS in any case.

VI. CONCLUSION

What is clear, is that applications based on such paradigms as we have described, and many envisioned for the future (and certainly not limited to the aerospace or space exploration domain) will be far too complex for humans to address all issues. Moreover, many issues will not be foreseeable, and much behavior will require hard real-time deadlines that can never be met with more traditional approaches.

While there is an overhead to achieving autonomicity, we believe that this overhead comes with significant benefit for RTS. We do not believe that it is too “costly” 1 for Real-Time Systems, but rather than in the future it will prove to be essential for developing effective RTS. Simultaneously, Autonomic Computing, and related areas, draw on much of the excellent research produced by the RTS community, a significant proportion of which was essential in making the AC initiative feasible.

In short: we believe we are moving swiftly towards a time when it will be imperative to have self-managing Real-Time Systems.

ACKNOWLEDGMENT

This work was supported in part by Science Foundation Ireland grant 03/CE2/I303_1 to Lero—the Irish Software Engineering Research Centre (www.lero.ie).

This research is partly supported at University of Ulster by the Computer Science Research Institute (CSRI) and the Centre for Software Process Technologies (CSPT) which is funded by Invest NI through the Centres of Excellence Programme, under the EU Peace II initiative.

Several of the technologies described in this paper were developed with NASA and are patent-pending and assigned to the United States Government.

REFERENCES

- [1] M. G. Hinchey and R. Sterritt, “Self-Managing Software,” *Computer*, 39(2):107-109, February 2006.
- [2] P. Horn, “Autonomic computing: IBM perspective on the state of information technology,” IBM T. J. Watson Labs, NY, 15th October 2001
- [3] R. Sterritt, M. G. Hinchey, “Why Computer-Based Systems Should be Autonomic,” *Proceedings of 12th Annual IEEE International Conference and Workshop on the Engineering of Computer Based Systems (ECBS 2005)*, Greenbelt, MD, USA, 3-8 April, 2005, Pages 406-414
- [4] J. O. Kephart and D. M. Chess. “The vision of autonomic computing”. *Computer*, 36(1):41–52, 2003.
- [5] R. Sterritt, “Autonomic Computing,” *Innovations in Systems and Software Engineering*, Vol. 1, No. 1, Springer, ISSN 1614-5046, Pages 79-88, 2005
- [6] R. Sterritt, M. G. Hinchey, “Birds of a Feather Session: “Autonomic Computing: Panacea or Poppycok?””, *Proceedings of 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, USA, 3-8 April, 2005, Pages 335-341
- [7] R. Sterritt, D. W. Bustard, “Autonomic Computing: a Means of Achieving Dependability?” In *Proceedings of IEEE International Conference on the Engineering of Computer Based Systems (ECBS'03)*, Huntsville, AL, USA, 7-11 April 2003, pp. 247-251.
- [8] R. Sterritt, “Pulse Monitoring: Extending the Health-check for the Autonomic GRID,” In *Proceedings of IEEE Workshop on Autonomic Computing Principles and Architectures (AUCOPA 2003) at INDIN 2003*, Banff, AB, Canada, 22-23 August 2003, pp. 433-440.
- [9] R. Sterritt, “Towards Autonomic Computing: Effective Event Management,” In *Proceedings of 27th Annual IEEE/NASA Software Engineering Workshop (SEW)*, Maryland, USA, 3-5 December 2002, IEEE Computer Society Press, pp. 40-47.
- [10] R. Sterritt, D. F. Bantz, “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, 2004.
- [11] R. Sterritt and D. W. Bustard, “Towards an Autonomic Computing Environment,” In *Proceedings of IEEE DEXA 2003 Workshops - 1st International Workshop on Autonomic Computing Systems*, Prague, Czech Republic, September 1-5, 2003, pp. 694-698.
- [12] R. Sterritt and M. G. Hinchey, “From Here to Autonomicity: Self-Managing Agents and the Biological Metaphors that Inspire Them,” *Proceedings of Integrated Design & Process Technology Symposium (IDPT 2005)*, Beijing, China, 13-17 June, pp 143-150.
- [13] S. M. Lewandowski, D. J. Van Hook, G. C. O’Leary, J. W. Haines, L. M. Rossey, “SARA: Survivable Autonomic Response Architecture,” In Proc. DARPA Information Survivability Conference and Exposition II, Vol. 1, pp. 77-88, June 2001.
- [14] W. F. Truszkowski, M. G. Hinchey, J. L. Rash, and C. A. Rouff, “Autonomous and Autonomic Systems: A Paradigm for Future Space Exploration Missions,” *IEEE Trans. on Systems, Man and Cybernetics, Part C*, 2006.
- [15] W. F. Truszkowski, J. L. Rash, C. A. Rouff, and M. G. Hinchey, “Asteroid Exploration with Autonomic Systems,” *Proceedings 11th IEEE International Conference on Engineering Computer-Based Systems (ECBS), Workshop on Engineering Autonomic Systems (EASe)*, Brno, Czech Republic, 24-27 May 2004, pp 484-489, IEEE Computer Society Press.
- [16] N. Kawasaki, “Parametric study of thermal and chemical nonequilibrium nozzle flow,” M.S. thesis, Dept. Electron. Eng., Osaka Univ., Osaka, Japan, 1993.
- [17] J. P. Wilkinson, “Nonlinear resonant circuit devices (Patent style),” U.S. Patent 3 624 12, July 16, 1990.
- [18] *IEEE Criteria for Class IE Electric Systems* (Standards style), IEEE Standard 308, 1969.
- [19] *Letter Symbols for Quantities*, ANSI Standard Y10.5-1968.
- [20] R. E. Haskell and C. T. Case, “Transient signal propagation in lossless isotropic plasmas (Report style),” USAF Cambridge Res. Lab., Cambridge, MA Rep. ARCRL-66-234 (II), 1994, vol. 2.