Towards Autonomic Networks

S. Schmid[†], M. Sifalakis[‡], D. Hutchison[‡]

[†] NEC Europe Ltd.
Network Laboratories
69115 Heidelberg, Germany
schmid@netlab.netlab.nec.de

[‡]Computing Department
Lancaster University,
Lancaster, LA1 4WA, U.K.
{mjs,dh}@comp.lancs.ac.uk

Abstract. Autonomic networking set a challenge for the research community to engineer systems and architectures that will increase the QoS and robustness of future network architectures. However, our experience is that so far the autonomic network research community does not have a common perception of what an *autonomic network* is. This paper attempts to propose a generic model for *autonomic systems*, along with a minimum set of required properties that would render a system compliant to this model. The paper emphasises the importance of such a common model for the credibility of the research community as well as to eliminate attempts to unnecessarily overload or blur the scope of the field.

Keywords: Autonomic communication, autonomic networks, autonomic system definition

1. Introduction

Over the last 15 years, the on-going convergence of networked infrastructures and services has changed the traditional view of the network from the simple wired interconnection of few manually administered homogeneous nodes, to a complex infrastructure encompassing a multitude of different technologies (wired/wireless, mobile/fixed, static/ad-hoc, etc.), heterogeneous nodes (regarding size, capabilities, power and resources constraints, etc.), diverse services (end-to-end, real-time, QoS, etc.), and competing objectives which are subject to "tussles-spaces" of interests [CWSB02]. This situation has put a challenge for the research community to engineer systems and architectures that will increase the QoS and robustness of the current and future internetwork whilst alleviating the management cost and operational complexity.

The autonomic communications research community [ACF] has been formed (among others) to respond to this challenge. Based on interdisciplinary grounds, it tries to tackle the problem by developing architectures and models of (networked) systems that can manage themselves in a reliable way always fulfilling their service mission. The need for such genuine course of action is also reflected in the memorandums of the initiatives that fund this research, in the EU [IST-FET] and world-wide (e.g. [NSF-FIND]).

Given the interdisciplinary grounds of the research area, it is very hard to argue about the originality of the practices, methods and theories that are anticipated by this work. There are a number of areas (see section 2) where their investigators would argue falls into the field of autonomic research. However, this paper emphasises that what characterises an area as novel/pioneering is the originality of its goals and objectives more than the means and methodologies used to achieve them. Besides, it is not surprising for different research communities to make use of the same or similar "tools" in order to achieve better solutions in their field. As a result, in order to improve the credibility of the autonomic research community as well as to eliminate attempts to unnecessarily overload or blur the scope of the field, this paper emphases the importance to have a clear and common view of what the research goals are and how to validate that. As a result, it is imperative that the research community agrees on what "autonomic" is and what it is not, and to have a way of accessing this functionality in a system.

Based on extensive discussion of this topic at various academic events [Dag06011][Infocom06][IWAN05][WAC05] and within the IST funded Autonomic Network Architecture (ANA) project [IST-ANA], this paper attempts to capture a definition for an autonomic system by proposing a model and identifying its essential properties, and advocates the need for a "test" that will enable the assessment of autonomic behaviour in a system. It is important to stress here that this paper does not claim to provide a complete definition, but it should rather be considered as a starting point for identifying the scope and the objectives of the research area.

Section 2 provides a short overview of early work in relevant areas, for example, autonomic computing. Section 3 provides some linguistic grounds to support the definitions that are propose in section 4. There it is also discussed what the essential properties of an autonomic system are. Section 5 advocates the need for a test to assess autonomicity in a system. Finally, section 6, concludes the papery with a summary of the work.

2. Background

Much of the research that has taken place in the past in various fields of computer science and network research could potentially classify as autonomic systems research (as argued by its investigators). Yet, as no standard definition exists to formally describe an autonomic system and its expected properties, the usage of the term remains vague and there is not conclusive way of assessing to what extend a research area qualifies as autonomic systems research and when the research objectives have been met. Nevertheless, there are several research areas that are relevant to the work on autonomic systems in a more or less direct way. This section points out a few of them.

Obviously and most probably the most relevant research work to autonomic systems seems to be the *Autonomic Computing* initiative of IBM through the e-Lisa project. Driven by their business market needs in 2001, IBM announced their intention to develop systems that are "autonomic". Inspired by the operation of the autonomic nervous system (ANS) in biological systems, they specified four properties that describe autonomicity in a system. These are:

- *Self-healing*: discover and repair potential problems to ensure that the system runs smoothly.
- *Self-protection*: identify threats quickly and take protective actions. Sensors feed data to a protection centre, for auditing, and action taking against various threats.
- *Self-configuration*: install and set up applications/patches/updates automatically, verify compliance with the specified service levels, optimise configuration of applications using adaptive algorithms.
- *Self-optimisation*: constantly monitor predefined system goals and performance levels to ensure that all systems are running at optimum levels.

In spite of the kudos created around this initiative and the seemingly visionary research effort by the sounds and the theory of it, IBM's vision proved to be rather short-sighted and focused on delivering simply software applications and application frameworks that could auto-configure, download updates, and occasionally learn user preferences or adapt to usage patterns with minimal administrative intervention, as opposed to infrastructures or architectures (software, computer systems, networks, etc) that can be self-managed and self-sustainable at software (system and application) and hardware level. Lest the term "computing" instead of "systems" in the definition.

Similar to IBM's effort to deliver "autonomic" business solutions are other projects by other market competitors such as Adaptive Enterprise [HPAE] by HP, Dynamic Systems Initiative [MSDSI] by Microsoft, eBusiness on Demand [IBMBOD] by IBM, N1 [SUNN1] by Sun Microsystems, Organic IT [FROIT] Forrester Research.

This paper differentiates between *autonomic systems* and IBM's *autonomic computing*. The reason being twofold: first, the term autonomic computing as defined by the IBM does not capture the complete spectrum of properties and behaviour an autonomic system should exhibit, and second by limiting our analysis to "systems" (whether that is a host, a network, or a complete infrastructure), the focus can be on well defined abstractions that are analysed, characterised and finally specified (section 4). This is in contrast to the term "computing" which has fuzzy and broad semantics and can include more or less anything and everything ranging from algorithms and software engineering methods to software and hardware systems.

Other areas related to autonomic systems include active and programmable networks [TW96, TSSM97]. The main connecting element between the two areas is the perceived need for adaptation and customisation of a system. An autonomic system must be able to customise the self in response to changes in policies or mission or in order to adapt to the deployment environment. Active and programmable platforms (e.g. [SFSS01, KRGP02, Hja00]) provide a means for generic adaptation and customisation. However active or programmable systems alone are not to be considered autonomic. Similarly to programmable systems, composable architectures (e.g. [HP91, MKJK99]) and component models (e.g. [CBGJLU04]) provide a degree of adaptiveness and customisability that is absolutely essential to autonomic systems.

Many people in the research community influenced by the exhibited autonomicity in biological systems tend to believe that an autonomic system needs to have intelligence and cognition. To that extend research on artificial intelligence and robotics

[Bro91] maybe closely related to autonomic systems. After all one of the long haunting visions of AI, has been to deliver robotic systems that think and act for themselves in the environment, i.e. autonomic systems per se. Evolutionary strategies [Sch01, Koz92] and probabilistic learning [Bis95] may offer sustainable solutions for adaptation in operational conditions that could not be determined, evaluated or prescribed at system design time.

While the authors fully support the argument that intelligence, learning and cognition can augment autonomic systems, this paper examines to what extend intelligence is an axiomatic property for autonomicity. On one hand experience has shown that deterministic systems and analytic methods may perform poorly in finding solutions for problems or conditions outside their initial design space therefore necessitating a degree of evolvability that draws from intelligence. On the other hand lessons learned from control theory show that it is possible to develop highly adaptive systems to unthought operational conditions without need for architectural evolvability. The main question to ask here is if one of the axiomatic properties of an autonomic system is the ability to evolve its architecture or simply adapt its operation in face of conditions that extend its original design space.

Finally another research area which often considered relevant to autonomic systems involves dependable and resilient systems (e.g. Grids). This is due to the inherent need of such infrastructures for persistence and survivability with low management effort and often in an unsupervised manner. An interesting point to notice here is that very often in these systems the building elements of are not autonomic individually but the whole system might be behaving autonomously collectively (from a high level management perspective). A misconception, which in our opinion is the differentiating point between the two research areas is the belief that an autonomic system is one which is self managed and "always ON", tolerant to attacks, problems, and condition changes. While it is true that an autonomic system must exhibit a reactive behaviour to counter problems, attacks and operational condition changes, it is not very clear what this reaction will entail. In some occasions it may be sensible for an autonomic system to hibernate during an "operational winter" period (e.g. after a natural disaster where the power may be a scarce resource) or even self-destroy depending on the mission (e.g. to avoid having secret codes intercepted or being hijacked). This behaviour might be contradicting with the principle of resilience or "always ON" operation. Although solutions are possible in all these cases our aim is to point out the potentially conflicting objectives of the two research areas which should not be identified as overlapping.

3. Some Etymology

A common misunderstanding comes from the fact that researchers believe autonomicity is the same as robustness, adaptability, intelligence and dependability, and as a result argue that research into autonomic systems is nothing new. However, this paper highlights that autonomicity, although related to the above topics, is something semantically different and new if considered holistically. The paper therefore tries to build clear understanding of what "autonomic" means and what the expected properties of an autonomic system are.

From a linguistic point of view the following definitions are relevant:

- *Automatic* [Greek automatos: auto- (*auto)* + -matos (*willing)*] refers to an act happening or an entity existing through the operation of a pre-existing arrangement (law of enforcement) that is triggered by some event.
- *Autonomous* [Greek autonomos : auto- (*auto)* + nomos (*law)*] means self-governed and refers to an entity capable of responding, reacting, or possibly developing independently of the whole, in mind or judgment and not controlled by others or by outside forces
- *Adaptive* means to posses the ability to change and make suitable to or fit for a specific use or condition.
- *Aware* means to have knowledge or experience of something and so being well informed of what is happening in that subject (state) at the present time. In accordance to this definition *self-awareness* means to have good knowledge and judgment about yourself

Although *autonomic* has the same roots as the word *autonomous* and is often used interchangeably (and consequently so advocated in IBM's terminology), from another linguistics point of view, autonomic lies between "autonomous" and the Greek word "autonomistic", which in turn refers to the tendency to be autonomous under any circumstance, functionally independent and under no voluntary control.

4. Defining Autonomicity

Although research in autonomic communications has its focus on network and communications aspects, the discussion regarding a definition for autonomicity aims to be more general and therefore tackles the problem from an abstract viewpoint. In this way, our definition tries to capture the semantics of autonomic systems in general, independent from whether it is software or hardware, an end system or a network system, etc.

From an architectural point of view, a *system* has a set of *inputs* and a set of *outputs*. A system implements some type of processing function, and based on its inputs, it generates some output. Zero inputs give the system semantics of a *source* and zero output semantics of a *sink.*

The following definition of an *autonomic system* stems from:

a) the semantics conveyed in the etymological study presented in section 2,
b) the empirical understanding of equivalent biological systems (ANS, eco-systems, living beings), and
c) the perception of the "common usage" of the term by the research community (i.e. considering those characteristics that are commonly associated with the term).

Proposed Definition:

An *Autonomic System* is a system that operates and serves its purpose by managing its own self without external intervention even in case of environmental changes. A conceptual model of an autonomic system is shown in (Figure 1).

A fundamental building block of an autonomic system is the sensing capability (*Sensors Si),* which enables the system to observe its external operational context. Inherent to an autonomic system is the knowledge of the *Purpose* (intention[1]) and the *Know-how* to operate itself (e.g. boot-strapping, configuration knowledge, interpretation of sensory data, etc.) without external intervention. The actual operation of the autonomic system is dictated by the *Logic,* which is responsible for making the right decisions to serve its *Purpose*, and influence by the observation of the operational context (based on the sensor input).

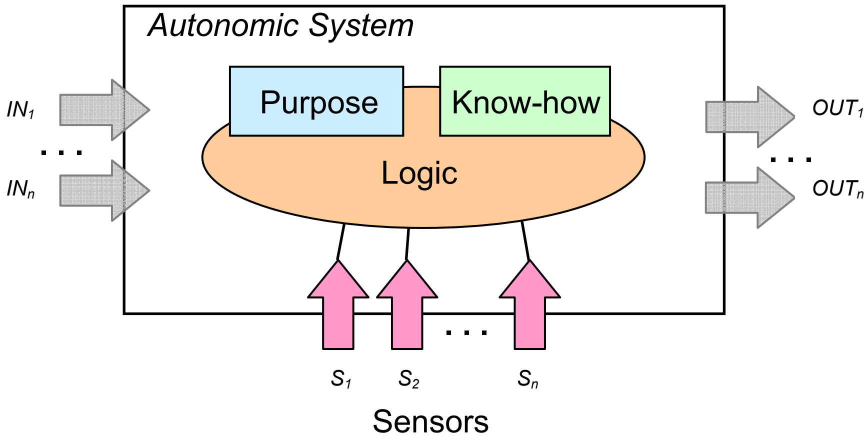

Fig. 1. Conceptual Autonomic System Model

This definition highlights the fact that the operation of an autonomic system is purpose-driven. This includes its mission (e.g. the service it is supposed to offer), the policies (e.g. that define the basic behaviour), and the "survival instinct". If seen as a control system this would be encoded as a feedback error function or in a heuristically assisted system as an algorithm combined with set of heuristics bounding its operational space.

Even though the purpose and thus the behaviour of autonomic systems vary from system to system, every autonomic system should be able to exhibit a minimum set of properties to achieve its purpose:

- **Automatic:**
  This essentially means for a system to be able to self-control its internal functions and operations. To be truly automatic (i.e. even to start-up the system) implies that a system is self-contained and able to bootstrap and operate without any manual intervention or external help (i.e. the *Know-how* required to bootstrap and operate the system must be inherent to the system).
- **Adaptive:**
  An autonomic system must be able to change its operation (i.e. its configuration, state and functions). This will allow the system to cope with temporal &

[1] Note that *intention* refers to the purpose or aim of an autonomic system.

spatial changes in its operational context either long term (environment customisation/optimisation) or short term (exceptional conditions such as malicious attacks, faults, etc.).

- **Aware:**
  An autonomic system must be able to monitor (sense) its operational context as well as its internal state in order to be able to assess if its current operation serves its purpose. Awareness will control adaptation of its operational behaviour in response to context or state changes.

Notice that the above definitions are still in-line with IBM's basic vision of *autonomic computing*. However, in contrast to the herein proposed definition, which defines the properties that characterise an autonomic system, IBM's four "autonomic *self*-properties", namely self-healing, self-configuration, self-optimisation and self-protection, define a set of functionalities or features that an autonomic system must provide. For example, according to our definition a system that can operate on its own while serving its purpose is autonomic, irrespective of whether it implements those functionalities. In our view, it should be the nature and purpose of an autonomic system that defines which functions are required! As a result, it can be argued that the herein proposed definition is more precise (in describing a system) and at the same time more general.

The above model is also consistent with the definition of a control system as well that of an AI agent. This is reasonable since practically a control system or an AI agent may easily implement an autonomic system, if it exhibits the aforementioned properties.

Finally, a short remark regarding the relation of autonomicity and evolvability is provided, since it has been often argued in discussions that an autonomic system ought to be evolvable (for example, through some type of artificial learning methods). Similarly to the above discussion regarding IBM's autonomic computing features, it can be argued that learning and evolvability may be a useful feature in an autonomic system, but whether it is required or not depends on the actual purpose of the autonomic system, and hence should not be considered an essential property of an autonomic system.

4.1 Why are these properties essential?

The question whether the three properties (automatic, adaptive, aware) are essential for autonomic systems or not requires some further discussion. This section follows a *reduction ad absurdum* approach to show that all three of them are essential**.**

First of all let's consider a system that is not automatic, yet adaptive and aware. Such a system has knowledge of its operational context and internal state, and can adapt current operation in response to context changes. However it cannot bootstrap on its own, or initiate functions that it has not used (needed) before, or terminate functions that it does not need anymore without external intervention. As a result, sooner or later it will not be able to serve its purpose anymore, either because it cannot start, or because it will not be able to perform a required function that it has not activated before (e.g. to counter an attack), or because it may run out of resources. Therefore such a system cannot qualify as autonomic.

If a system is automatic and aware but not adaptive, then it can bootstrap alright and control its possible operations, as well as sense the environment and know its internal state. However, if the environment changes, although the change will be detectable, the system will not be able to decide what steps it needs to take in order to adjust accordingly its operation with regard to serving its purpose. The ability to sense the environment and monitor internal operation (awareness) and to control the operation (automatic) is not necessarily sufficient for the system to be able to service its purpose. Operational changes involving internal (re-)configuration and/or state changes will not be feasible if a system lacks the ability to adapt.

Finally let's assume a system which is automatic and adaptive but not aware. Following a similar line of argumentation it is easy to infer that such a system will not be autonomic because it will not be able to trigger adaptation and eventually not survive in face of environment changes. Most important, such a system has no way of sensing and assessing if its operation complies with its *purpose*.

As a result of this analysis, it can be concluded that all three properties are essential for a system to be autonomic.

4.2 Autonomicity and Deterministic Behaviour

The question whether an autonomic system needs to behave/operate in a deterministic way has often been raised. For an autonomic system to be deterministic would imply that its behaviour will be fully predictable, and thus its compliance with its purpose could be verified.

According to our definition, deterministic behaviour is not one of the essential properties of an autonomic system. The reason lies in the conjecture that in any non-analytical or non-combinatorial system, although it is not possible to determine or predict the "next state change" or "action", the system can still be engineered in such a way that it will comply with its purpose (using bounds or heuristics). This is also true for biological systems, where often in similar situations different decisions are made, but still to achieve the same objective (e.g. hunt to feed).

Deterministic behaviour and autonomicity should therefore be considered orthogonal concepts. Whether or not an autonomic system is deterministic depends on the logic that defines the system behaviour, which in turn depends on the type of system and its purpose. For example, if the logic of an autonomic system is able to self-evolve through learning, the behaviour of the system easily becomes non-deterministic. On the contrary, if the logic is defined by a fixed finite state machine, the systems behaviour is deterministic and therefore traceable.

In general it may be desirable to have a deterministic autonomic system, as one can easier verify its operation and behaviour. On the other hand, this may also limit an autonomic system in a way. For example, it doesn't allow the system to learn from its "experiences" and self-evolve.

5. Testing Autonomicity

Up to now, a model for understanding the basic functionality and properties that are required by a system to exhibit an autonomic behaviour has been proposed. However, when implementing an autonomic system one needs to be able to verify and assess to what extend the goal is met and if the modelled functionality is reflected in the operation of the system.

This is a particularly difficult problem to address, which can be partly attributed to the lack of empirical experience with real platforms or simulation models of autonomic systems due to the young life of the research field.

One possible approach to address this problem is by testing one by one the four properties that characterise an autonomic system. Separate simple tests could be developed that could be used to test the behaviour of a system in different scenarios and verify to what extended the system is automatic, adaptable and aware. Then a more synthetic test should ascertain the soundness of the expected interactions (involving the *logic*) between these four properties in fulfilling the *purpose* of the system.

Another approach would be to seek for a more abstract "Turing-like" test whereby the system is seen as a black box, with the only verifiable being what one can perceive externally: (a) *is the system is still operational?* (b) *is the expected service is still provisioned?* However, like in the original Turing test there is no possible way of testing it against all possible cases.

A last approach is to continuously test the system at run-time (on-line test). A test function can be periodically applied in the system to assess its compliance with the *purpose*. In this case the challenge is to determine a set of metrics to "measure" automatically the behaviour of the system. These metrics would need to be developed on a case per case (system) basis.

Testing evolvable and thus potentially non-deterministic autonomic systems is even more challenging, since in this case it is not possible to simulate the *logic* and test the system off-line (only on-line tests may be applicable). Practically, there is no obvious way of validating the correctness of the autonomic behaviour of a self-evolved system (as it is unclear how the systems changes its operational pattern), and therefore it may not even be possible to identify a self-evolving autonomic system.

6. Conclusions

The main objective of this paper is to identify and demonstrate the need for means of accessing to what extend a system exhibits an autonomic behaviour. First, the paper proposes and analyses a model for an *autonomic system*. Then, a set of essential properties that render a system compliant to this model have been identified. Finally, a set of possible options for testing these properties in order to inspect the autonomicity of a system have been discussed. The incentive for this lies in the argument that unless there is a common understanding of what is expected from an autonomic system and how to assess autonomicity, there is no common ground to drive research in the area of autonomic communication and networking. Such a foundation is also needed to establish a credible community with well defined objectives, tools and methodologies.

This is of fundamental importance to sustain longer term research and evaluate its success.

Acknowledgments. This work has originated within the ANA project which is an EU-IST funded project under the FET initiative. We would also like to thank many of the participants of the Schloss Dagstuhl seminar 06011 on *"Autonomic Networks"* for inspiring and motivating this work.

References

[ACF] Autonomic Communications Forum. Online at: http://www.autonomic-communication.org/

[Bis95] C. M. Bishop. "Neural Networks for Pattern Recognition". Oxford University Press (1995).

[Bro91] Brooks, R.A. "How to build complete creatures rather than isolated cognitive simulators," in K. VanLehn (ed.), Architectures for Intelligence, pp. 225-239, Lawrence Erlbaum Assosiates, Hillsdale, NJ, 1991.

[CBGJLU04] Coulson, G., Blair, G.S., Grace, P., Joolia, A., Lee, K., Ueyama, J., "A Component Model for Building Systems Software", In *Proc. IASTED Software Engineering and Applications* (SEA'04), Cambridge, MA, USA, Nov 2004.

[CWSB02] Clark, D. D., Wroclawski, J., Sollins, K. R., and Braden, R. 2002. Tussle in cyberspace: defining tomorrow's internet. In *Proc. of the 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications* (SIGCOMM 2002), Pittsburgh, Pennsylvania, USA, 2002.

[Dag06011] International Conference and Research Center for Computer Science, Schloss Dagstuhl, Wadern, Saarland, Germany, January 2006.

[FROIT] Organic IT. Forrester Research. Online available at: http://www.forrester.com/ Research/Document/ Excerpt/ 0,7211,34342,00.html

[Hja00] Hjalmtysson, The Pronto platform - a flexible toolkit for programming networks using a commodity operating system. In Proc. of the International Conference on Open Architectures and Network Programming (OPENARCH 2000), March 2000.

[HP91] N. C. Hutchinson and L. L. Peterson, "The x-kernel: An Architecture for Implementing Network Protocols," IEEETransactions on Software Engineering, vol. 17, January 1991.

[HPAE] Adaptive Enterprise Services and Adaptive Network Architecture. HP Invent. Online available at: http://h20219.www2.hp.com/services/cache/77602-0-0-0-121. html

[IBMBOD] E-Business On Demand. IBM Research. Online available at: http://www-306.ibm.com/e-business/ ondemand/us/index.html

[Infocom06] 25th Conference on Computer Communications (IEEE Infocom), Barcelona, Spain, April 2006.

[IST-FET] Future and Emerging Technologies Program. Information Society Technologies. Online available at: http://cordis.europa.eu/ist/fet/home.html

[IST-ANA] AutonomicNetwork Architecture Project. Online available at: http://www.ana-project.org

[IWAN05] 7th International Working Conference on Active Networks (IFIP IWAN), Sofia Antipollis, Cote d' Azour, France, November 2005.

[Koz92] Koza, J.R., Genetic Programming: On the Programming of Computers by Means of Natural Selection. MIT Press 1992.

[KRGP02] Keller, R., Ruf, L., Guindehi, A., Plattner, B., PromethOS: A Dynamically Extensible Router Architecture Supporting Explicit Routing. In Proc. of the 4th International Conference on Active Networks (IWAN), Zurich, Switzerland, December 4-6, 2002.

[MKJK99]R. Morris, E. Kohler, J. Jannotti, and M. F. Kaashoek. The Click modular router. In Proc. of the 17th ACM Symposium on Operating Systems Principles (SOSP '99), pages 217--231, Kiawah Island, South Carolina, December 1999.

[MSDSI] Dynamic Systems Initiative. Microsoft Corporation. http://www.microsoft.com/windowsserversystem/dsi/default.mspx

[NSF-FIND] Future Internet Network Design Initiative. National Science Foundation. Online available at: http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=12765

[Sch01] Schmitt, Lothar M (2001), Theory of Genetic Algorithms. Theoretical Computer Science (259), pp. 1-61.

[SUNN1] N Computers Operating as 1. Sun Microsystems. Online available at: http://www.sun.com/software/ learnabout/n1/

[SFSS01] Schmid, S., Finney, J., Scott, A., C., Shepherd, W., D., Component-based Active Network Architecture, In Proc. of IEEE Symposium on Computers and Communications, July 2001.

[TW96] D. L. Tennenhouse and D. J. Wetherall. Towards an active network architecture. In Proc. of Multimedia Computing and Networking '96, January 1996.

[TSSM97] DL. Tennenhouse, JM. Smith, W. D Sincoskie, DJ. Wetherall, and G J. Minden. A Survey of Active Network Research. *IEEE Communications Magazine*, Vol. 35, No. 1, pp80-86. January 1997.

[WAC05] 2nd International Workshop on Autonomic Communication (WAC IFIP 2005). Vouliagmeni, Athens, Greece Oct. 2005.