



A Role for Contextualised Knowledge in Autonomic Communications

Sterritt, R., Mulvenna, M., & Lawrynowicz, A. (2004). A Role for Contextualised Knowledge in Autonomic Communications. In *Unknown Host Publication* Springer.

[Link to publication record in Ulster University Research Portal](#)

Published in:
Unknown Host Publication

Publication Status:
Published (in print/issue): 01/10/2004

Document Version
Publisher's PDF, also known as Version of record

General rights
Copyright for the publications made accessible via Ulster University's Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy
The Research Portal is Ulster University's institutional repository that provides access to Ulster's research outputs. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact pure-support@ulster.ac.uk.

A Role for Contextualised Knowledge in Autonomic Communications

Roy Sterritt, Maurice Mulvenna, Agnieszka Lawrynowicz

¹School of Computing and Mathematics, University of Ulster at Jordanstown,
Newtownabbey, County Antrim, BT37 0QB Northern Ireland
r.sterritt@ulster.ac.uk md.mulvenna@ulster.ac.uk

²Institute of Computing Science, Poznan University of Technology,
ul. Piotrowo 3a, 60-965 Poznan, Poland
agnieszka.lawrynowicz@cs.put.poznan.pl

The conceptual architecture of autonomic communications requires a knowledge layer to offer effective, transparent and high level self-management capabilities. This knowledge plane can utilise the behaviour of autonomic communication regimes to monitor and intervene at many differing levels of network granularity. This paper introduces autonomic computing and autonomic communication, before outlining the role of knowledge in autonomic networks. Some research issues, in particular the concept of dynamic context as a method to acquire knowledge dynamically that helps to facilitate a successful realisation of the knowledge plane are explored and discussed.

Introduction

An EU FET brainstorming workshop in July 2003 to discuss novel communication paradigms for 2020 identified '*Autonomic Communications*' as an important area for future research and development [1]. This can be interpreted as further work on self-organizing networks, but is undoubtedly a reflection of the growing influence of IBM's Autonomic Computing initiative launched in 2001 [2]. In effect, autonomic communications has the same motivators as the autonomic computing concept with particular focus on the communications research and development community. Goals highlighted at this initial workshop were to understand how an autonomic network element's behaviours are learned, influenced or changed, and how in turn, these effect other elements, groups and networks. The ability to adapt the behaviour of the elements was considered particularly important in relation to drastic changes in the environment such as technical developments or new economic models [1].

At the heart of autonomic communications are *selfware* principles and technologies that will create the autonomic network. They borrow largely from autonomous distributed systems research and non-conventional networking (ad hoc, sensor, peer-to-peer, group communications, active networks and so on), among others [3]. In addition to this a new construct, a knowledge plane, has been identified as being

required to act as a pervasive system within the network that builds and maintains high level models of what the network is supposed to do in order to provide the communications services and advice to other elements in the network [4]. It is generally considered that this knowledge plane will rely on the tools of AI and cognitive systems (to meet the uncertainties and complexities of this goal) rather than traditional algorithmic approaches [4][5].

This paper motivates the proposition that the successful creation of autonomic communications, and in particular the knowledge plane, requires the ability to possess context awareness and behavioural knowledge from an ethnomethodological perspective. Ethnomethodology is an in-depth study of individuals and groups, their practice, and their artefacts in the context of their normal working environment. From this perspective, context is more than just the sum or function of the metrics that are monitored or probed in the environment.

The paper outlines the area of autonomic computing and autonomic communications before beginning to discuss the role of knowledge in autonomic communications. The remainder of this vision paper is a discussion on how knowledge may be used within the knowledge plane of autonomic communications. In particular, we examine the mechanism of dynamic context as a framework for the generation, use and execution of knowledge in autonomic networks.

Autonomic Computing

The autonomic metaphor, based on the human body's autonomic or self-regulating and protection system, strives to achieve systems which will maintain themselves through the use of an autonomic element consisting of an autonomic manager and the managed component. There is a strong requirement for dependability, from single mobile devices running multiple processes through distributed grid applications [6].

The general properties of an autonomic (self-managing) system can be summarised by four objectives; self-configuring, self-healing, self-optimising and self-protecting and four attributes; self-awareness, environment-awareness, self-monitoring and self-adjusting [6]. Essentially, the objectives represent broad system requirements while the attributes identify basic implementation mechanisms. (Since the 2001 launch of autonomic computing the *self*-* list has grown substantially yet this initial set still represents the general goal.)

Self-configuring is a system's ability to readjust itself automatically, this may simply be in support of to changing circumstances or to assist in self-healing, self-optimisation or self-protection. Self-healing, a reactive mechanism is concerned with ensuring effective recovery when a fault occurs; identifying the fault and then where possible repair it. Self-optimisation means that a system is aware of its ideal performance, can measure its current performance against that ideal and has policies 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 a mix of self-configuration and self-optimisation to increase protection. Finally, these self-mechanisms should ensure there is minimal disruption

to users, avoiding significant delays in processing. To achieve these objectives a system must be aware of its internal state (self-aware) and current external operating conditions (environment-aware). 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 of how these may be adjusted. The ability to operate in a heterogeneous environment requires the use of open standards to understand and communicate with other systems [1].

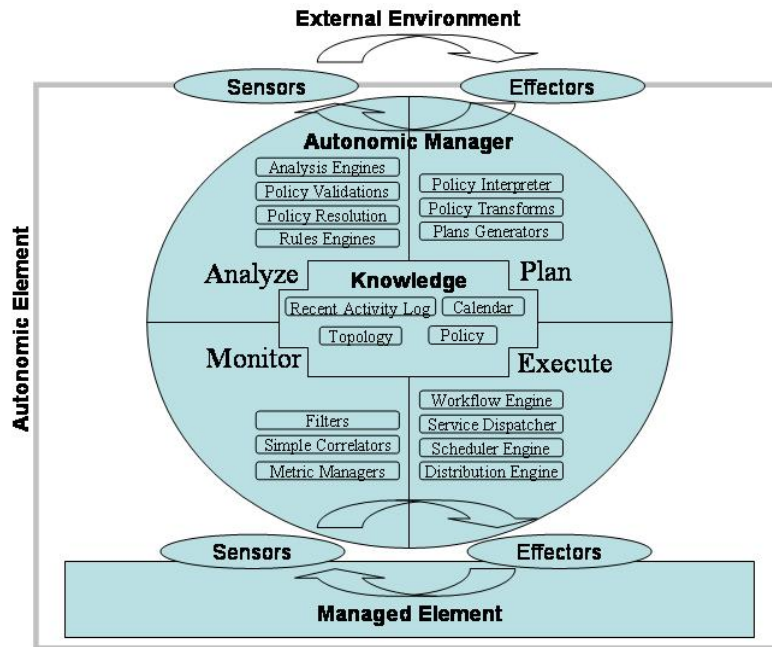


Fig. 1. IBM’s view of the Architecture and Components of an Autonomic Element (adapted from [7])

At the heart of any autonomic system architecture are *sensors* and *effectors* [8]. A control loop is created by monitoring behaviour through sensors, comparing this with expectations (historical and current data, rules and beliefs), planning what action is necessary (if any) and then executing that action through effectors [9]. The control loop, a success of manufacturing science for many years, provides the basic backbone structure for each system component [7].

Figure 1 is IBM’s view of the necessary components within an autonomic manager. (For an alternative artefacts view, see [10].) It is assumed that an autonomic manager is responsible for a managed element within a self-contained autonomic element. Interaction will occur with remote autonomic managers through virtual, peer-to-peer, client-server [11] or grid [12] configurations.

The monitor and analyse 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 behaviour that will be executed through the effectors. The *MAPE* (Monitor, Analyse, Plan and Execute) 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 autonomic manager (AM).

Autonomic Communication

The traditional perspective for autonomic computing is that an autonomic element (AE) solely *uses* knowledge; there is no explicit *creator* or *adaptors* of knowledge. It implies that the knowledge within is *engineered in* as part of the developed autonomic manager (and updated from an external source). If you consider the management scope and assume this autonomic manager's component is for instance a disk drive, engineering the knowledge may be achievable yet if the scope is larger, for instance a higher level manager within a server farm receiving event communications from many other autonomic managers, the scenarios will be too complex to engineer.

AI may assist here. When autonomic computing first hit the headlines it was interpreted as being AI or at least containing it, perhaps due to the commonality of the biological analogies. IBM has spent effort stating this is not the case to thus avoid the initiative being classed with the past perceived failures of AI hype. Here lies a difference between autonomic computing and autonomic communications, as autonomic communications accepts the intrinsic need for AI to create the knowledge plane [4].

In this autonomic computing view, even if you do assume that AI and machine learning techniques have been used to assist in developing/engineering the rules and beliefs, another question arises as to how adaptable these are within the autonomic manager.

Proponents of the mobile and/or intelligent agent paradigm would present that context drives adaptability through agent's capability to discover, extract, interpret and validate context [13], and as such will enable them to make a significant contribution to the autonomic communications field. This is not in doubt; a wide range of techniques will be required for the successful creation of autonomic communications. Yet will emergent behaviour from autonomic elements agents provide the scope envisaged at the knowledge plane level?

An interesting paper in [14] discusses affect and machine design [15]. Essentially it supports those psychologists and AI researchers that hold the view that affect is essential for intelligent behaviour. It proposes three levels for the design of systems:

1. Reaction – lowest level where no learning occurs but immediate response to state information coming from sensory systems.
2. Routine – middle level where largely routine evaluation and planning behaviours take place. It receives input from sensors as well as from the reaction level and

reflection level. This level of assessment results in three dimensions of affect and emotion values: positive affect, negative affect and (energetic) arousal.

3. Reflection – top level receives no sensory input or has no motor output, it receives input from below. Reflection is considered a meta-process, where the mind deliberates about itself. Essentially operations at this level look at the systems representations of its experiences, its current behaviour, its current environment etc.




Fig. 2. Intelligent Machine Design three tiers compared with three planes in Autonomic Communications

The affect and emotion debate is not an issue here, it is the three levels that are of specific interest. Although not described in such terms this approach to intelligent design is similar to the proposed scoping of the planes within autonomic communications (Figure 2). Essentially the reaction level may be considered to sit within the data plane and the autonomic network, where for instance under fault conditions automated switching and fail-over may take place, and so on in, monitoring current state of both the network element and its environment with rapid reaction to changing circumstances. The routine level may be considered the management plane, where planning takes place and under fault conditions root cause analysis is performed on the event messages from the data plane. The reflection level may be considered not to reside yet within networks but is akin to the perceived knowledge plane, where it considers the behaviour of the networks and learns new strategies, reflects upon the success of existing strategies and adapts if necessary.

This approach highlights the need for reflection and cognitive strategies to be designed into systems to provide the self-adaptability autonomic property.

Self-adapting behaviour has been classified into three levels by the Smart Adaptive Systems community. These are [16]:

1. Adaptation to a changing environment,
2. Adaptation to a similar setting without explicitly being ported to it,
3. Adaptation to a new/unknown application.

It seems a difficult task for an autonomic element and indeed the autonomic network and autonomic communications to even conform to level 1 through engineering rules into the autonomic manager. To be classified level 2 is certainly going to entail AI and cognitive approaches.

This section has briefly focused on the initial general designs of an autonomic element to emerge from autonomic computing, key criteria consisting of self-managing (self-CHOP¹ or self-*), AE=AM+ME, MAPE, control loop, sensors+effectors, all reliant on knowledge repository to provide self and environment awareness. These *users* of the knowledge component have yet to identify how this knowledge will be learnt, adapted or even used within different contexts.

The Role of Knowledge

Behavioural knowledge and knowledge execution is a vital research area for the successful fulfilment of Autonomic Communication [17]. In order to drive the self-managing capabilities of autonomic communications, there is a requirement for the network to be self-aware and environment-aware. Research on self-awareness in next-generation networks can be driven by attempting to understand the behaviour of the network. To achieve this, the network must have access to various data and knowledge components, on which it can execute and modify its parameters.

The data and knowledge sources are [18]:

- Deriving and using first- and second-order data from the data plane of the network;
- Deriving and using network management data and knowledge from the control plane of the network;
- Deriving and using data and knowledge that comprises the knowledge plane [4] of a network.

The first two data and knowledge sources can be and are employed to varying degrees in network research today. The third area represents a significant advance in research thinking, in that it is primarily inferential and mined knowledge that is discovered by predictive analytic techniques residing on the knowledge plane. These techniques include collaborative filtering [19][20], Bayesian networks [21], clustering [22], classification [23], association rules [24], sequence analysis [25] and content filtering [26] as well as runtime techniques from click stream analysis [27][28].

The knowledge plane must have the capacity to retain and maintain a *network memory*, comprising the data and knowledge sources indicated above. An excellent starting point for this memory will be machine-understandable XML-based syntax, comprising different standards that maintain high semantic integrity and coherence for the knowledge models; for example, the Predictive Modelling Mark-up Language (PMML) [29].

PMML is an XML-based standard developed by the Data Mining Group with the aim of aiding model exchange between different model producers and between model producers and consumers. Most data mining vendors have their own proprietary representations for knowledge discovered using their algorithms. PMML provides the first standard representation that is adhered to by all the major data mining vendors. Being XML-based, models represented in PMML can easily be parsed, manipulated and used by automated tools. The anticipated use of flexible, semantically-rich representational schemes such as PMML within autonomic elements is as a memory

¹ Objectives: self-Configuring, self-Healing, self-Optimising and self-Protecting

for policies and events that provides fast, interactional response in autonomic network environments.

This network memory will be maintained as a discrete ontological construct in the knowledge plane, necessitating new research in network ontologies. This memory is, in essence, a collection of rule sets and mining model result sets that can maintain network policies as well as behavioural descriptions and policies. As such, it is a memory that provides context for measurement. Therefore, via introspection and mediation, the memory can self-adapt to improve performance depending on the context and needs of use.

In order to execute and interact with the network memory, a scalable high-performance engine is required. This is similar in construct to a recommender engine [27][28], in that it is constantly updating the network memory rule bases upon which the application of predictive algorithms on network behavioural data is based.

A key component of this engine is the detection of network trends and subtle changes in data flows. Key research currently under way in concept drift may be the basis for drift detection in autonomic network architectures [30].

There are key challenges in this research sub-area of autonomic communications, including the real time handling and assessment of ensembles of behavioural knowledge to improve network provision and the ability to introspectively measure the performance, accuracy and appropriateness of network performance.

Contextualised Knowledge

The autonomic communications knowledge plane not only requires the ability to *use* knowledge but also the ability to *create* and *adapt* it when necessary. A vital aspect to these abilities is to understand the *context* within which that knowledge is framed.

The understanding of *context* has been a significant research area in many fields of computing, in particular AI and ubiquitous computing, for some time now. The term context-aware computing was first introduced by in 1994 [31] as a system's ability to adapt to its location of use and objects (people, devices) in the neighbourhood. It was defined in the context of the systems in which the user employs many different mobile, embedded and stationary computers in different situations and locations over the course of the day. This has evolved within several research fields sharing many common views, including ubiquitous computing [32], pervasive computing [33], ambient intelligence [34], planetary/utility/grid computing and so on.

Many definitions of context-awareness and models of context-aware systems have been proposed, the most popular over-arching perspective that researchers from pervasive computing society employ is to see context as some function or mode of the parameters of the environment, such as time, place, etc. Values of the parameters are acquired by using the predefined set of sensors and then extract features from these low-level sensor readings [35].

Acquiring context is not a straightforward task due to its dynamic nature and the heterogeneous state of data sources. Context can be extracted from low-level sensors and high level managers as well as derived from applications to-date utilising the network. It has been highlighted that the majority of context-aware applications use

the data from the sensors later offline through data pre-processing and features extraction [35].

There is no consensus for context representation (capturing, representing and modelling context). Problems concern the fact that acquired information can be strongly heterogeneous and often incorrect, inconsistent or incomplete. A second issue is that it is used in systems in various ways. A substantial amount of different approaches have been proposed to model such contextual information.

Dourish [36] has suggested that the representational approach of context applied by most of the researchers until now interprets the role of context in a different manner than it plays in our everyday life. He proposes instead a new perspective on context-aware computing where the context is perceived much like in social sciences that study the practises of individuals in their normal environment. In his article he examines the problem of context from a high-level, philosophical point of view, enumerates the various philosophical viewpoints, and highlights an approach, which views context as an interactional problem rather than a representational one.

This *dynamic context*, as we term it, contrasts with a majority of the literature concerning context-awareness, particularly engineering approaches that inherit from positivist theoretical tradition which seek objective answers, independent of the detail of particular occasion descriptions of social phenomena. It is a positivist point of view in which we look at things as something to be modelled and encoded. From this perspective, context is a stable feature of the world that is independent of the actions of individuals.

On the other hand, dynamic context proposes to look at the problem of context-aware computing from another, *phenomenological* point of view. In this view, social facts are not pre-given or absolute but are continually negotiated and reinterpreted as a result of interactions hence perception of the world depends on the interpretation of particular individuals.

In dynamic context, it is the activity that generates and sustains the context. So, context arises from the activity, and is actively produced and maintained in the course of the activity. This provides a framework for a method to determine context from activity via behaviour (and measures of behaviour). This framework is a justifiable research goal in autonomic communications.

Although the assumptions enumerated above seem to be a correct way to view the role that context should play in context-aware systems, there are many significant issues concerning how to turn this approach into reality. Dynamic context is only a conceptual view of what context is and formal design guidelines for systems are not presented. It is an interactional model of context, in which the central problem is “how and why, in the course of their interactions, do people achieve and maintain a mutual understanding of the context for their actions?” It can be argued that the difficulty and practical problems of designing context-aware systems has encouraged the pervasiveness of the representational view of context.

Context is hard to recognise and hard to encode. The approach of dynamic context makes this task even harder, because instead of the readings from the set of predefined sensors we have to deal with the features that can be contextually relevant to the particular activity. It cannot be determined *a priori*, before an activity happens. Some features become meaningful for particular sorts of actions - that’s why the context should be continually redefined, as such the scope of contextual features

should be defined dynamically. That forces the representation to be flexible enough to maintain the changing importance of the features in different types of activities and their dependence together with the possibility to add or delete features. Dourish gives a conceptual idea of how context should be understood and suggests to move the stress from designing how to use the predefined context within a system, but rather how the system can support the process by which “context is continually manifest, defined, negotiated, and shared”.

A Framework for Using Context

In dynamic context, the activity generates and sustains the context. This fresh perspective is drawn ultimately from social science techniques such as ethnomethodology and ethnography, and explores their usefulness in the increasing number of computer-mediated pervasive and ubiquitous environments. Ethnomethodology simply means the study of the ways in which people make sense of their social world [37]. Ethnography is the in-depth study of individuals and groups, their practices, and their artefacts in the context of their normal work environment [38].

The usefulness of ethnography seems to be that it takes nothing for granted, and the application in anthropology ensures that all details are available for analysis. Ethnography is a contender for a framework in which we seek to discover context from activity. Having the measures of behaviour, we can then try to discover from them the activity and from the activity we can try to discover the context. This leads to the exciting unexplored possibility for a new general framework for context-aware computing.




Fig. 3. Interactional Vs Representational view of context-awareness

Having a set of sensors within the autonomic elements, each sensor provides the measures of network, systems and user application behaviour. On the basis of the AE’s measurements we can try to discover the application’s current activity. Frequent patterns found in the measurements can then be labelled and represent typical activities. Having discovered activity we can try to compute the context generated by the given activity, understand a more precise description of the activity in the form of

its goal, and the conditions in which it is executed (that means context viewed by the prism of activity). The general Behaviour-Activity-Context (BAC) Framework for context-awareness is shown in on the left-hand side of Figure 3.

Discussion and Proposed Research Plan

The knowledge plane is a proposed third abstraction in the emerging research area of autonomic communications, adding to the existing data and control/management planes. In their vision paper, the proponents of the knowledge plane discuss broadly how machine learning algorithms can be applied to garner knowledge and increase the self-awareness of the network. How the knowledge plane will be achieved is an open research area, but the remaining discussion examines what role contextualised knowledge may play in autonomic communications.

The paper first focused on a brief review of the general autonomic element designs emerging from autonomic computing noting that the general architecture of the autonomic computing autonomic element would imply it is only a *user* of knowledge with no explicit components for *creating* nor *adapting* knowledge. Agents, AI and cognitive techniques may assist here. It was highlighted that the three tiers; reflex, routine and reflection of the proposed autonomic intelligent machine design have scope commonalities with the data, management and knowledge planes within autonomic communications.

The second EU FET consultation meeting in March 2004 on the subject of autonomic communications [39][40][41] highlighted that self-awareness in autonomic communications must be driven by self-knowledge, specifically by behavioural knowledge. This key area was entitled behavioural knowledge and knowledge execution. The authors put forward several areas of knowledge research which they feel should be pursued to support the use of behavioural knowledge in autonomic communications:

- The use of unsupervised, incremental learning algorithms should be explored. Although there are many machine learning and data mining algorithms available, comparatively few researchers have explored this area, in particular from a pervasive computing perspective [36].
- The second of these is the development and use of existing research and tools that facilitate high-performance operation; specifically, ontological tools to support the incorporation and use of semantic information.
- Knowledge systems in autonomic communications should be capable of practicing introspection. That is, they should measure the degree of correctness of their 'advice' within an autonomic element.
- The knowledge system should be capable of discriminating between conflicting types of advice and selecting or blending advice. This can be explored initially as simple conflict resolution, but a key goal would be the development of managers of ensemble advisors or recommenders within autonomic communications.

More generally, this paper has introduced the concept of dynamic context and advanced the proposition that the successful creation of autonomic communications

and the knowledge plane will not only require AI and cognitive approaches but will also require a fuller interpretation of context; in some ways akin to ethnography, building towards the formulation of a novel context-awareness framework; Behaviour-Activity-Context (BAC).

A treatment of the area of context-awareness highlighted the two schools of thought; *interactional* versus *representational*, or *phenomenological* versus *positivist* perspective. This paper supports the interactional view where context is generated and sustained by the activity. This ethnographical-inspired view of the world, which we have labelled dynamic context, should provide the most dynamic knowledge approach for autonomic communications.

This paper explicitly focused on one of the grounding principles to achieve autonomic communications – a new communication paradigm to assist the design of the Next Generation Networks (NGN) – that of contextualised knowledge.

We have proposed a new dynamic context model, based upon on-the-fly, dynamic and lightweight analysis of data in the network, as well as a workable framework for experiments. We propose a research plan that tests the hypothesis that contextualised knowledge can improve the capabilities of a knowledge plane in autonomic communications. The details of the roadmap for this plan will be explored in a future paper.

References

1. EU IST FET, "New Communication Paradigms for 2020", brain storming meeting July 2003, Brussels, Belgium, (report published Sept 2003)
2. P. Horn, "Autonomic computing: IBM perspective on the state of information technology", AGENDA'01, October 2001
3. M Smirnov, R Popescu-Zeletin, "Autonomic Communication", presentation EU IST FET brainstorming meeting Communication Paradigms for 2020, Brussels, July 2003
4. D Clark, C Partridge, JC Ramming, JT Wroclawski, "A Knowledge Plane for the Internet", Proc. Applications, technologies, architectures, and protocols for computer communication, Karlsruhe, ACM SIGCOMM 2003
5. JM Agosta, S Crosby, "Network integrity by inference in distributed systems", NIPS Workshop on Robust Communication Dynamics in Complex Networks, 2003
6. R Sterritt, DW Bustard, "Autonomic Computing - a Means of Achieving Dependability?", Proc. IEEE Int. Conf. Engineering of Computer Based Systems (ECBS'03), Huntsville, AL, USA, , pp 247-25, April 7-11 2003
7. A Ganek, "Autonomic Computing: Implementing the Vision", Keynote presentation at the Autonomic Computing Workshop, (AMS 2003), Seattle, WA, 25th June 2003.
8. AG Ganek TA Corbi, The dawning of the autonomic computing era, IBM Sys J 42(1) 5-18 2003
9. Autonomic Computing Concepts, White Paper, IBM, 2001
10. R. Sterritt, DW Bustard, "Towards an Autonomic Computing Environment", 1st Int. Workshop Autonomic Computing Systems at DEXA'2003 Prague, 694-698 2003.
11. DF Bantz, C Bisdikian, D Challener, JP Karidis, S Matrianni, A Mohindra, DG Shea, M Vanover, Autonomic personal computing, IBM Sys J 42(1) 165-176 2003
12. G. Deen, T. Lehman, J. Kaufman, The Almaden OptimalGrid Project, IEEE Proc. Autonomic Computing Workshop, (5th AMS), Seattle, WA, pp 14-21, June 2003.

13. A. Zaslavsky, Mobile Agents: Can they assist with Context Awareness?, Proc. 2004 IEEE Int. Conf. Mobile Data Management, 2004
14. IBM Systems Journal, Special issue on Autonomic Computing, Vol. 42, No.1, 2003
15. DA Norman, A Ortony, DM Russell, "Affect and machine design: Lessons for the development of autonomous machines", IBM Sys J, 42(1), 38- 44, 2003.
16. D Anguita, "Smart Adaptive Systems: State of the Art and Future Direction of Research", EUNITE, Dec. 2001
17. M Smirnov, 'Area: Autonomic Communications', EU IST FET New Communication Paradigms for 2020 Consultation meeting, Brussels, Belgium. (ver. 02), March 2004.
18. MD Mulvenna, Annex: Comments and background material to topic 8 Behaviour knowledge and knowledge execution in Autonomic Communications, IST FET New Communication Paradigms for 2020 Consultation meeting, Brussels, March 2004.
19. U Shardanand, P Maes, Social information filtering: Algorithms for automating "word of mouth", Proc CHI'95 Human Factors in Computing Systems, 210-217 1995
20. P Resnick, N Iacovou, M Suchak, P Bergstrom, J Riedl "GroupLens: An open architecture for collaborative filtering of netnews", Proc ACM 1994 Conf Computer Supported Cooperative Work, Chapel Hill, NC: ACM, 175-186 1994
21. D Heckerman, D Geiger, D Chickering Learning Bayesian networks The Combination of Knowledge and Statistical Data, Machine Learning, 20 197-243 1995
22. AK Jain, RC Dubes, Algorithms for Clustering Data. Pren Hall, NJ 1998
23. D Mitchie, DJ Spiegelhalter, CC Taylor, Machine Learning, Neural and Statistical Classification, 1994. www.amsta.leeds.ac.uk/~charles/statlog/
24. R Agrawal R Srikant, Fast Algorithms for Mining Association Rules. In Proc. of the 20th Int'l Conference on Very Large Databases, Santiago, Chile, September 1994.
25. AG Büchner, M Baumgarten, SS Anand, MD Mulvenna, JG Hughes, Navigation Pattern Discovery from Internet Data, Advances in Web Usage Analysis and User Profiling, LNCS, Springer-Verlag, 2000.
26. M Mobasher, R Cooley, J Srivastava, Automatic Personalization Based On Web Usage Mining, Communication of ACM, August, 2000, Volume 43, Issue 8
27. MD Mulvenna, SS Anand, AG Büchner, (eds.), Personalization on the Net using Web Mining, Comm. ACM Special Section, 43(8), 122–125, Aug. 2000
28. AG Büchner, MD Mulvenna, Discovering Marketing Intelligence Through Online Analytical Web Usage Mining, ACM SIGMOD Record, 27(4), 54-61, 1998
29. PMML, www.dmg.org
30. M Black, R Hickey, Learning classification rules for telecom customer call data under concept drift. Soft Comput. 8(2): 102-108 (2003)
31. B Schilit, N Adams, R Want, R. Context-Aware Computing Applications. Proc. IEEE Workshop on Mobile Computing Systems and Applications, 1994.
32. M Weiser, The Computer for the 21st Century. Sci American, 265(3), 94-104 1991.
33. W Ark, T Selker, A Look at Human Interaction with Pervasive Computers, IBM Sys J, 38(4), 504-507 1999
34. E Aarts, R Collier, E van Loenen, Bd Ruyter, (Eds.), Ambient Intelligence, 1st European Symposium, EUSAI 2003, Veldhoven, The Netherlands, LNCS 2875 2003
35. R. Mayrhofer, H. Radi, et al. Recognizing and predicting context by learning from user behaviour, Austrian Computer Society (OCG), 2003.
36. P. Dourish, What we talk about when we talk about context. Personal and Ubiquitous Computing 8(1) 19-30, 2004
37. Garfinkel, H., *Studies in Ethnomethodology*, Prentice-Hall, 1967
38. Bowling, A. (1997). *Measuring health: A review of quality of life measurement scales* (2nd ed.). Philadelphia: Open University Press.

39. M Smirnov, Managing Internet complexity in Autonomic Communication, presentation EU IST FET consultation meeting Communication Paradigms for 2020, Brussels, March 2004
40. EU IST FET, "New Communication Paradigms for 2020", Consultation meeting 3-4th March 2004, Brussels, Belgium.
41. F Sestini, 'Situated and Autonomic Communications', EU IST FET New Communication Paradigms for 2020 Consultation meeting, Brussels, March 2004.