Panel Report: "Main principles to guide R&D in Algorithms, Protocols and Middleware"

Ioannis Stavrakakis[1], Fabrizio Sestini[2]

[1] University of Athens ioannis@di.uoa.gr, , [2] EU Commission, Fabrizio.Sestini@cec.eu.int

Abstract: the panel objective was to discuss with the audience and highlight commonalities and inter-dependencies between the papers presented in the first two sessions of WAC 2004 - **Network Management** and **Models and Protocols**. As this panel was the first one, main discussion tried to define autonomicity, address the more general questions around this new concept and distinguish it from previously introduced and studied concepts.

What is autonomicity?

Do we need a new formalism, unifying abstraction or new paradigm to describe and manage it?

There was some consensus on the need for a scalable policy language in a highly distributed-based networking scenario, whereas references to agent technology or to active and intelligent networking should be avoided. However fundamental aspects that lead to new networking paradigms need to be effectively addressed first, such as the behavior and cooperation of autonomic entities (based on considerations of economics and game-theoretic approaches) and the self-organizational aspects (and associated effective algorithms).

The issue is not simply the individual autonomic behavior serving self-interest, but how these autonomic, self-driven, individual behaviors lead to a desirable, acceptable global behavior (this is addressed in biology in the context of behavioral ecology). In this respect feedback from the network and control aspects are also important.

We should also clearly articulate to what extent autonomic communications is different from re/auto-configurability, which should be considered as just one aspect of it.

Why do we need autonomic communications?

Internet is indeed complex, but is working pretty well and satisfying a lot of requirements. The following three reasons in favour of development of autonomic communication were expressed during the panel.

First, to cope with the increasing complexity of the Internet by enabling autonomicity, i.e. new networking solutions that facilitate injecting an autonomic component into the network and figuring out automatically what to do with it.

Second, to share the cost of managing networks (this comes from a similar reasoning as to why autonomic computing is needed – IBM's desire not to be burdened with the cost of mainframes).

Third, the cell size is decreasing in mobile wireless networks and such networks proliferate. It is therefore highly desirable to deploy and manage them autonomously and not manually, to reduce the increasing management cost.

Is autonomicity more of a revolutionary than an evolutionary concept?

To claim that this research is more revolutionary as opposed to evolutionary, it was suggested not to concentrate only on the current or emerging problems it may help solve, but rather try to establish if it fits a possible trend, by looking fairly far into the past of this technology and trying to foresee its future.

60 years ago networks had one or very few owners (in a country or a continent) and this worked well under the circuit-switching paradigm (which did not see any revolution over those years). 30 years ago, the revolutionary packet-switched Internet technology was formed and is managed now by a very large number of distributed entities. Today, we have end users being capable of providing network resources (as routers, storage as content proxies) to networks and becoming themselves the network. Thus, it is conceivable that, several years in the future; the network will basically be built, owned and run by the millions of autonomic end-users. Getting to and managing such networks requires revolutionary steps, which can justify why we need autonomic communications. The trend suggests that it is very likely that autonomic nodes will shape the networking environment in the future, contributing their increasingly powerful resources to it.

Consequently, in addition to developing the concept of autonomicity in order to re-design networking to be more effective in the current and emerging networking environment, we should develop this concept in such a revolutionary way as to capture early enough emerging trends that will lead to entirely new networking environments where autonomic communications would clearly be the paradigm to employ.

Panel Report: "Grand Challenges of Network and Service Composition"

Giuseppe Valetto[1], Fabrice Saffre[2]

[1]Telecom Italia Lab, Giuseppe.Valetto@tilab.com,
[2]British Telecom, fabrice.saffre@bt.com

Abstract: This brief report intends to summarize some of the things we learned during WAC2004 sessions "Network Composition" and "Negotiation and Deployment", and that were highlighted during the subsequent panel discussion. We try to focus especially on the following aspects: traits and trends of convergence emerging from the rather diverse findings presented in those sessions; controversial or divergent opinions on some of those findings; open issues that should be addressed by the autonomic communication community and how to tackle them; and major research directions that seem likely to emerge and shape a significant part of the autonomic communication landscape.

Findings and Commonalities

We notice how a certain set of issues surfaced repeatedly in both sessions. Figure 1 attempts to draw some relationships among those issues, starting from the idea of *Adaptation* as a generic term for the self-regulating operation of an autonomic system. Several of the works presented deal with *Composition* as an important form of adaptation and a first-class autonomic primitive. *Policies* and *Negotiation* are widely regarded as useful means to strategize about and achieve adaptation through composition, in a top-down vs. bottom-up fashion: policies can dictate the terms of composition (e.g. *when, what and how*), while negotiation can be used to spontaneously reconcile competing and/or conflicting policies, and converge towards stable system configurations. Finally, *Semantic Knowledge* can be used to obtain an explicit and abstract representation of autonomic adaptation, including composition. That representation remains formal and hence can be distilled into policies, but at the same time provides the ability to understand discuss, communicate and review the autonomic behavior of a system.

Another evident common element is that wireless communication provides major motivation and an obvious test bench for the investigation of autonomic paradigms: LANs, PANs, sensor networks, as well as ubiquitous, ambient and ad hoc networking, are among the wireless and mobile contexts that look like natural catalysts for autonomic capabilities. It seems also evident that all of them require some kind of context-awareness (e.g. location-, user-, service-awareness) embedded in the autonomic communication facilities, possibly as part of the knowledge base.

Other agreed-upon, significant enablers of autonomic communication are transparent addressing, seamless handover, strong decentralization of all adaptation mechanisms, and a regard for issues like resource, security and trust management as first-class elements.

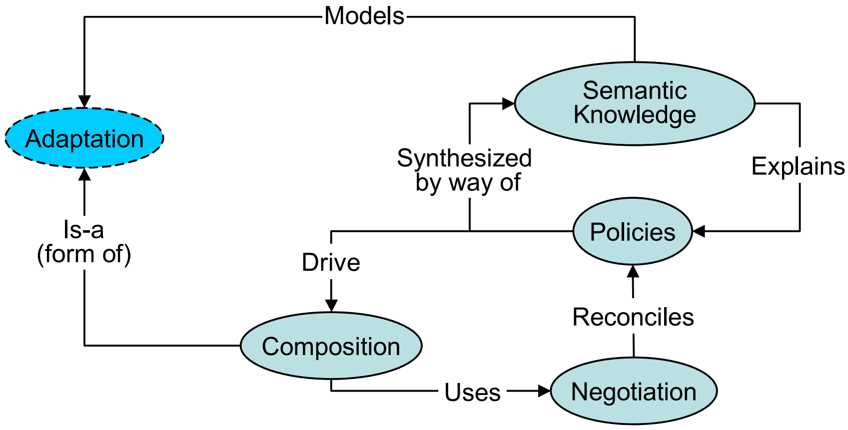

Figure 1: Common issues and their relationships.

Divergences

Along many of the presentations and discussions, the issue of *emergent behavior* was often either explicitly mentioned, or alluded to. However, we have recorded diverse opinions on its actual relevance to the problems in autonomic communication. On the one hand, it seems clear that it has a clear appeal, because of its resonance with the biologic metaphor at the very basis of all things autonomic, and its promise to enable complex, numerous and strongly autonomous elements that co-exist in the same environment to act in concordance and with a common sense of purpose. On the other hand, there is a feeling that the biologic metaphor should not be over-stretched, and a concern that such an approach – as well as many evolutionary or "cognitive" approaches - could break once applied to communication infrastructures and services, under their extremely demanding timeliness and predictability requirements.

In a similar fashion, different opinions exist on whether autonomic features should originate from explicit vs. implicit provisions (e.g., dedicated protocols that account for and codify features such as robustness, flexibility and fault-tolerance vs. spontaneous or stigmergetic interactions among communication elements that converge towards a mutually sustainable and satisfactorily functional configuration).

Open Issues and Next Steps

A number of suggestions and questions have been raised, trying to indicate how to focus the undergoing work and discussion of the autonomic communication community towards a set of incremental goals that can help map this largely uncharted territory.

A major issue seem to be trying to define the boundaries of the autonomic communication domain: it seems obvious that it comprises self-* issues within a network infrastructure, as well as at the juncture of different networks; it is also evident that it extends to a degree up to the level of the services carried by those networks, but what kind of services are going to be affected the most – and in what application domains - is not equally well understood.

A related issue is the categorization of the techniques that are relevant for autonomic communication: in Session 2A and 2B – and throughout the Workshop in general – a wide spectrum of techniques, ranging from control theory (with its hard mathematical foundations), to bio-inspired techniques (for instance, emergent behavior), and from distributed negotiation algorithms to knowledge-based reasoning have been presented and debated. It must be better understood what solutions are suitable for what problems, in a multi-dimensional space that includes qualitative parameters, such as problem type, attainable scale, level of impact, as well as quantitative properties, such as timeliness, predictability, reliability.

Another open issue of great interest of the community is the level of *transparency* that autonomic communication facilities shall strive for. Transparency is in itself a multi-faceted concept, since it encompasses technical factors (e.g., non-intrusiveness in affected systems) as well as human-observable factors (e.g., the ability to understand, analyze and influence autonomic behavior on the part of technicians, as well as seamless and "hidden" operation from the point of view of end users).

An operational suggestion to address some of the questions above is to work together to propose and formalize a sort of "autonomic communication reference problem" (or problem set), against which proposed approaches should be evaluated, independently of their application domain. The reference problem might be characterized as a "check-list" of observable and demonstrable autonomic features, which an autonomic communication system should strive to address, thus enabling a form of comparison with respect to other solutions.

Research Directions

Research issues in autonomic communication are likely to revolve around achieving two complementary objectives: identifying suitable design principles and testing implementation/deployment strategies. As for the first, many sources of inspiration are available from the natural world including, but not limited to, biological systems. As for the second, a number of well-established domains in computer science (machine learning, formal methods etc.) offer reliable tools and a substantial body of knowledge to start with. This is not the place to review these, however, so we chose to sketch a picture of what we believe will be a fundamental common denominator to all future research in the field of autonomic systems.

Traditional engineering starts by specifying desirable system-wide characteristics and then designs/selects individual components under the assumption that the whole is only the sum of its parts. In extremely large distributed systems, this “top-down” approach is under severe strain to deliver viable solutions, which is a major rationale for autonomic communication. Yet this new paradigm raises issues of its own, mostly due to the apparition of complexity (in the restrictive scientific sense of emergent global properties within large ensembles of interacting units). However, despite being aware of this difficulty, many technologists seem reluctant to cross the cultural barrier between a proven and immensely successful paradigm (inherited from the industrial revolution) and the newer science of complexity, which is less well understood by engineers.

Complexity science provides powerful methods for dealing with probabilistic predictability and describing in a rigorous and useful way systems comprised of individually unpredictable elements. Over the last three decades, it has been extensively demonstrated that variability in the individual response of its constituents does not necessarily translate into the frequency distribution of a system's states exhibiting a similar amount of 'noise'. On the contrary, the huge number of interactions and the presence of intricate feedback loops often mean that the system as a whole can only exist in a limited number of configurations, despite the largely random behavior of individual units. The science of complexity mainly consists of identifying these configurations, determining their probability of occurrence, and understanding/characterizing transitions between them and trajectories leading to them (e.g. bifurcation).

The sheer size of a large network comprised of many thousands of components means that the state of a large distributed computing environment will virtually always be the result of an unforeseeable combination of many events, and so can be described best probabilistically. While there may be an increased recognition of this situation, there is a poor awareness of the methods capable of dealing with it. The heterogeneity of the underlying infrastructure (in terms of purpose, capability, and ownership) precludes a centrally imposed set of rules defining the function and privileges of every participant. Instead, we must find ways to engineer autonomic principles, like self-configuration, into individual elements and their interactions, so as to allow them to deal with unexpected situations, requests, combinations of events, etc.

Complex systems theory and modelling can and must help us understand which macroscopic behaviour is more or less likely to emerge from the many interactions between heterogeneous devices. The real challenge is not to cope with microscopic unpredictability - the conceptual tools required to handle its macroscopic effects are readily available. The difficulty resides in identifying and weighting the factors involved, so that the purpose of fine-tuning the local rules is not defeated by the presence of 'hidden variables' capable of pushing the entire system into an unexpected/undesirable state.

Panel Report: "How the Autonomic Network Interacts with the Knowledge Plane?"

David Lewis

Knowledge and Data Engineering Group,
Trinity College Dublin
Dublin 2, Ireland
Dave.Lewis@cs.tcd.ie

Abstract: This panel was held at the end of the Workshop on Autonomic Communication Principles, on the 19th October 2004. It brought together speakers from session 3 on Resilience and Immunity and session 4 on Meaning, Context and Situated Behaviour. The panellist were Anuarg Garg (University of Trento), Fabio Massacci (University of Trento), Christian Tschudin (University of Basel), Simon Dobson (University College Dublin), Maurice Mulvenna (University of Ulster) and Cesar Santivanez (BBN Technology).

Panel Report

The panel opened with a question from the audience asking how the evolvability of Autonomic Communications (AC) and the stability of the resulting architectures and systems can be ensured. The panel responded stating that stability can not be regarded in terms of deterministic system configuration, but needs to be viewed in terms of behavioural stability. Thus we must tolerate a level of volatility but only within a well understood behavioural envelope that relates to specific autonomic tasks. In other words, we should focus on enforcing specific bounds on the adaptivity that self-managed systems may exhibit, rather than on achieving full behavioural determinism. With respect to the evolution toward and evolvability of AC, it was agreed that gradual changes were a real-life necessity. As a result we require ways to subdivide AC architectures into separate areas of concern that can be attacked, solved and deployed independently. However, there were no immediate suggestions for the lines along which such a separation would best be made. Preceding the panel, a poster presentation had included a synthesis of issues raised during the workshop in the form of a layered cube reminiscent of that used to explain broadband ISDN principles during the 1990's. It was observed that this synthesis served to show the potential complexity and inter-connectiveness of issues in AC. Reactions to this model, however, also hinted at the challenges in defining any clear architectural separations for AC given our current understanding of the field. It also spurred comments on the lessons that could be learnt by the failure of ATM to reach its technical potential due to a lack of flexibly in reacting to changing economic and market concerns. There was broad consensus that these lessons must be heeded by the AC community in considering any evolution strategy.

Next a speaker from the audience observed that the success of the Internet was due in no small part to the clear separation of the application from the network via a simple interface, but that this separation also potentially limited the evolution of communication services. The question then posed asked whether the application-network separation should be subject to some re-integration to open the door to fundamental reappraisal of architectures suitable for autonomic communication. Such re-integration is already a strong feature in much current research in cross layer optimisation for wireless and ad hoc networks. Some panellists viewed that the application/network separation should not be violated due to its significance in allowing application innovation. Another emphasised the need for some form of modularity in order to allow the problem to be broken down and for innovation and competition to be encouraged. A further response questioned the assumption that the knowledge needed for AC should not be allocated to a separate 'plane' and suggested that instead it should be integrated with the data plane of the network. It seemed increasingly apparent from these responses that a layered architecture with well defined interfaces between layers was not readily apparent for AC. Instead, it was observed that the focus should shift to the adaptive sharing of information across conventional network boundaries, but in a way that was constrained by business, regulatory, or task concerns, rather than the need to have a fixed interface in support of a stratified architecture. However, as a result, the computational elements that populate such a loosely structured AC architecture must be more able to deal with information exchanged with other elements without pre-programmed understanding of its semantics.

The next speaker from the floor reinforced this view by observing the use of terms such as 'network of workflows' and 'architecture as a program' in the workshop. This was followed by a specific question on how AC systems can best determine when 'text becomes context', i.e., how is available knowledge to be judged useful context for a problem? Here the panel was broadly agreed that there is no canonical model of what represents context for AC tasks. Instead, context had to be formed on a subjective basis by AC elements, resolving their knowledge needs against the information that is available and accessible to them. This raises the prospect that the process that identifies and uses information as context determination is itself context-aware.

The final speaker from the floor asked how AC systems could be made conceptually simple. There was consensus from the panel that the problem domain was implicitly complex, and that the target should be to simplify the human experience of the management of complex communication services. It was observed however that we should not aim for one-size-fits-all approach to exposing complexity to the human administrator, but to instead aim for complexity on demand to reflect the tasks, skills, and cognitive abilities of the individuals concerned. However, though the complexity that individual AC components expose could be minimised, this is likely to be at the expense of increasing complexity in how such components interact.

The panel ended with each panellist expressing how what they had learnt in the workshop would impact on how they might subsequently present their papers. Christian Tschudin, whose paper presented a fine-grained approach to integrating code fragments that resulted in robustness and self-healing properties of the overall program, had his views on the need for a bottom-up approach reinforced. Fabio Massacci, who had presented a paper on negotiating the knowledge exchange needed to resolve access control policies, was interested in applying such a fine-grained approach to achieving robustness in policy integration. Simon Dobson, whose paper examined the role of contextual semantics in AC, would address the role of composition more carefully in the determination of the semantics of context. Maurice Mulvenna, whose paper had addressed the customisation context knowledge to the task at hand, was interested in the need for a more rigorous experimental approach to AC development. Anuarg Garg, whose paper had addressed a peer-to-peer trust mechanism combining concepts of reputation, quality, and credibility, expressed a need to more clearly define the relationship between P2P and AC. Cesar Santivanez, whose paper addressed adaptable ad hoc networks, saw the need to make ad hoc networks more application aware.

Conclusions

In conclusion, the role of a 'knowledge plane' in relation to Autonomic Communication remains unclear, in no small part due to a lack of consensus on what characterises such a plane. Clark et al's 2003 SIGCOM [1] paper described a 'Knowledge Plane for the Internet' as operating in parallel to existing concepts of data, control, and management planes. However, a closer examination of this work reveals that it encompasses not only knowledge monitoring and analysis but also its use for the planning and execution of network control and management tasks, thus making it much closer in functional scope to Autonomic Communications. Their use of the 'knowledge plane' metaphor probably owes more to the pragmatic tendency in the Internet Community to progress through a set of small, individually motivated steps rather than as part of a larger cohesive vision. Though the panel recognised the need for incremental evolution of AC, the aim of the workshop was to start work on a comprehensive AC vision and on the research agenda needed to realise it. As such, we are justified in questioning the core separation of layers and planes underlying the design of current networks, and in particular the persistence of this mindset into architecture for AC. The panel underlined this critical stance, raising the prospect that the AC domain may not be amenable to decomposition into the type of orthogonal separations that has guided the separation of concerns in current networks. This has profound implications for the AC research agenda and the resulting market in AC systems. Though an alternative architectural structure is not yet evident, some themes have been hinted at in the panel. These include the need: for composition of AC elements; for mechanisms to bound the adaptive behaviour of such compositions, and for mapping this adaptive behaviour to bounds on the behaviour of elements. Also raised is the need to tailor both the exposure of complexity and the employment of contextual knowledge, to the specific task at hand.

References

1. Clark, D., Partridge. C., Ramming, J.C., Wroclawski, J.T. “A Knowledge Plane for the Internet”, in Proc. of SIGCOMM’03, 25-29 August 2003, Karlsruhe, Germany