

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).

1 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 flexibility 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. Anurag 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.

2 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