Xi'an Jiaotong-Liverpool University 西交利物浦大学

2nd International Workshop on Smart Grid Technology and Data Processing
18 March 2016 in Suzhou, China

Smart networks for smart grids and smart cities

New network science: RINA and ΔQ

Peter Thompson
Chief Technical Officer, Predictable Network Solutions Limited, Bristol, U.K.
Peter.Thompson@pnsol.com

Abstract—**Smart grids, smart cities and the Internet of Things all require connectivity that is substantially more secure, resilient, predictable, scalable and efficient than packet networks are today. New approaches are required: here we consider the Recursive InterNetwork Architecture (RINA) and the quality attenuation calculus (ΔQ).**

Keywords-smart grid; smart city; IoT; RINA; quality attenuation; ΔQ

I. Introduction

The urgent problem of human-induced climate change, together with other problems such as particulate air pollution, is driving a rapid transition to renewable forms of energy generation [1]. Some sources of renewable energy, such as wind and solar photovoltaic, are intermittent. Compensating for this using only fossil fuel 'stand-by' generation reduces the environmental benefit and increases costs. Large-scale energy storage is one solution, but this is also expensive. An additional difficulty with renewable energy generation is that it is frequently geographically distributed, in contrast with the established model of centralised generation, transmission and distribution, causing strain on the system designed for the older model. These factors make it important to balance supply and demand on both a local and an aggregate basis, using 'smart grids'.

Smart grids arise from a convergence of: electricity generation, transmission and storage; distributed data processing; and embedded computation ('smart devices'). Away from the specific context of energy management, the last two technologies constitute what is called the 'Internet of Things' (IoT). When applied in a municipal context this can deliver 'smart cities', in which various problems such as traffic control, waste management, air pollution etc. are alleviated by combining information from arrays of sensors with models of urban behaviour, to control traffic lights etc. across the city [2].

From distributed computing perspective, smart grids, IoT and smart cities have similar requirements. In this paper we will consider these requirements, and how they can be met using new developments in network science, in particular the Recursive InterNetwork Architecture (RINA) and the understanding of quality attenuation (ΔQ).

II. Requirements for 'smart' networks

Existing network technologies have evolved from 1970s experiments in connecting general-purpose computers together to allow remote terminal access and file transfer; these computers were few in number, and each one was individually managed by a system administrator. The ARPANET extended this connectivity to a wide area, but within a largely closed and well-intentioned user group [3]. ISPs developed in order to provide public access to this system via dial-up modems, supporting email, file transfer, and http web access. With the advent of DSL technology, the public could have much higher speed and always-connected access, which enabled an explosion of new uses for the network, transforming commerce and society. The simplicity of a flat address space and best-effort delivery provide a low barrier to entry, enabling innovation; unfortunately this low barrier also applies to cybercrime/terrorism. Moreover, many uses of the network now critically depend on performance characteristics that are accidental consequences of design and implementation choices. Maintaining those characteristics in a best-effort context is becoming increasingly difficult and expensive, undermining the broadband delivery model.

Extending the same approach to the IoT and smart grids/cities is even more problematic; design choices made in the 1970s in response to a specific set of circumstances are now very far from optimal. We must now consider the interconnection of potentially very large numbers of devices that must operate reliably and unattended, in an environment that includes malicious actors. Instances of attacks on internet-connected infrastructure are becoming increasingly common, and a hack of a smart grid or smart city network could disable an entire region. Moreover the performance of the distributed computations is no longer just a matter of profit, convenience or entertainment. For example, the response of a smart grid to fluctuations of supply and demand is time-critical, in order to avoid brown-outs for consumers or even damage to the electrical transmission and distribution system.

We can formulate the requirements for networks to become part of the civic infrastructure as follows:

A. Security

The network should be 'secure-by-design', rather than requiring complex additional mechanisms to patch up inherent weaknesses. It should be straightforward to apply formal reasoning and proof tools to demonstrate security – or at least to show where proof obligations are missing and hence where security holes may lie. To do this at reasonable cost requires a uniform and composable structure, so that reasoning can be applied to subsystems and the conclusions combined to produce valid statements about the whole system.

B. Resilience

The network should be tolerant of node/link failures (which are inevitable as the scale increases). Failures should be contained and recovered from locally rather than propagating state changes across the network; otherwise the whole network can become unstable.

C. Predictability

The network should deliver predictable performance for given levels of demand and resource constraints, so that distributed applications can perform correctly with a high degree of certainty.

D. Scalability

The complexity of managing the network, and resources such as routing tables, should scale sub-linearly with its size. Managing complexity requires appropriate composable abstract models, so that management tasks can be decomposed into simpler tasks that do not depend on knowledge of, nor interaction with, the entire network configuration and state.

E. Efficiency

High-speed network interfaces inherently consume more power than lower-speed ones. For applications such as smart grids with very many small-scale nodes, it is essential for the capacity of network links to be used efficiently so that the interfaces can be as low-power as possible.

III. New Technology for Smart Networks

Having presented a challenging set of requirements, we now turn to new approaches that make meeting these requirements considerably easier. We will consider RINA, the Recursive InterNetwork Architecture, and ΔQ, a performance calculus for distributed applications.

A. RINA

The Recursive InterNetwork Architecture (RINA) is a new framework for communications that starts from the view that all communication is inter-process communication (IPC). Starting afresh from this perspective, a basic set of core capabilities have been defined [4]. As the name suggests, the architecture is uniform and recursive, in contrast to the fixed set of specialised network layers as found in OSI or TCP/IP. Specialisation occurs by means of 'policies' that enable each layer to optimally perform its appropriate set of

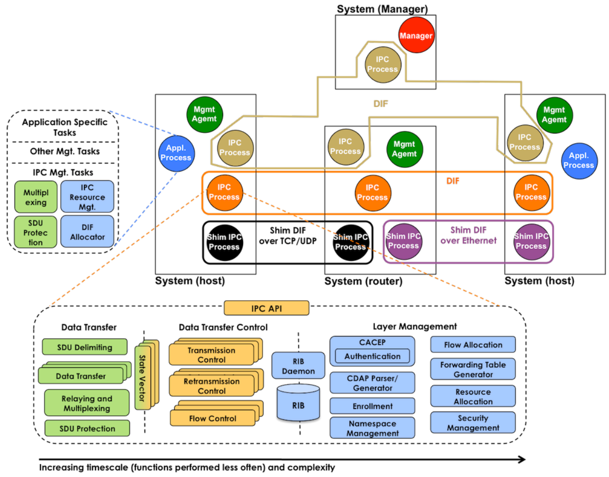

Figure 1: Components of RINA Architecture

functions. Instead of a single, flat, address space as in in IP, RINA provides a hierarchy of 'Distributed IPC Facilities' (DIFs), each with its own addressing scope.

This has three immediate benefits: it keeps address spaces small, improving scalability; it restricts the scope affected by a link or node failure; and it distinguishes higher-level identifiers relating to application processes from lower-level ones such as interface names, enabling seamless multi-homing and mobility. RINA is thus inherently more scalable and resilient than current approaches.

RINA includes a well-specified security model including clear trust relationships between different layers. RINA's nested addressing scopes also constrain or even eliminate typical data-transport attack approaches in IP networks such as port scanning and connection-opening [5]. DIFs are securable, providing an inbuilt distributed firewall capability. Moreover, the homogenous and recursive structure of RINA would support application of formal analysis and proof techniques at reasonable cost.

RINA is a relatively new approach, and still at quite an early stage of development. There are several on-going projects in the US and in Europe, co-ordinated by The Pouzin Society. One notable European effort is the FP7-funded PRISTINE project, which is close to delivering on its goals: to design, develop and implement the innovative internals of this architecture; and to demonstrate the applicability and benefits of this approach and its built-in functions in use-cases driven by the end-users, service providers and equipment vendors in the consortium [6].

B. ΔQ

The task of networking is often thought of as 'moving' a volume of data from one location to another. However, consonant with the RINA view of networking as IPC, we can focus instead on whether a set of communicating processes can achieve a given outcome within a bounded time. In any layered architecture, a higher-level outcome will depend on a sequence of lower-level ones (the sequence being determined by a protocol); the lowest-level 'outcome' may be simply

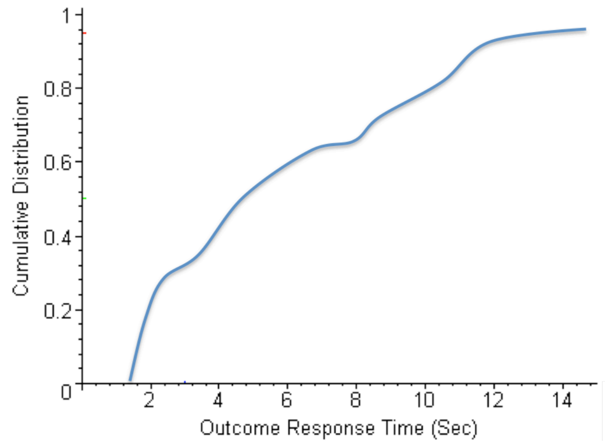

Figure 2: ΔQ as an Improper Random Variable

translocating the contents of a packet from one physical location to another. Given that high-level outcomes (having an acceptable voice conversation, watching an online movie, etc.) often depend on a very large number of lower-level ones, it makes no sense to consider every outcome individually; rather we must consider them statistically. In order to determine the probability of achieving a higher-level outcome, it is not in general sufficient simply to know the *rate* of the lower-level outcomes; what is needed is the probability distribution of how long they may take, and the chance they may fail altogether, which can be encoded by an improper random variable, as shown in **Figure 2**. This represents the deviation from perfectly instantaneous and infallible outcomes, 'attenuated' by the constraints of the real world, so it is called 'quality attenuation' and written 'ΔQ'.

Given the ΔQ of the lower-level outcomes, and the details of the protocol that turns sequences of them into a higher-level one, it is possible to calculate the ΔQ of the higher-level outcome, for example using stochastic process algebras [7] or Monte-Carlo simulations. **Figure 3** is a contour plot showing the interaction of delivered network quality degradation (the peak of a uniformly distributed delay and Bernoulli loss rate) with the user visible outcome – the time to complete the delivery of 10Kb of data in response to a HTTP (web page) request.

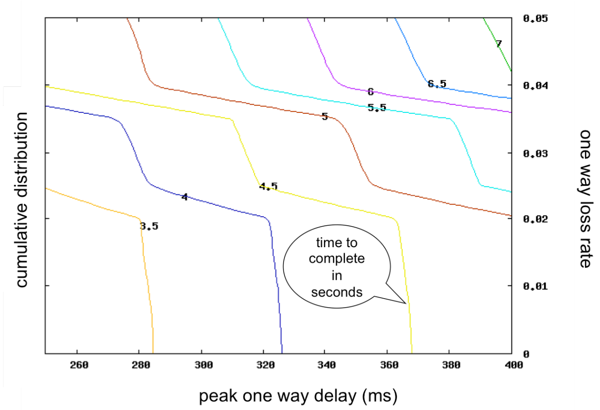

Figure 3: Median time to complete for 10kB HTTP request

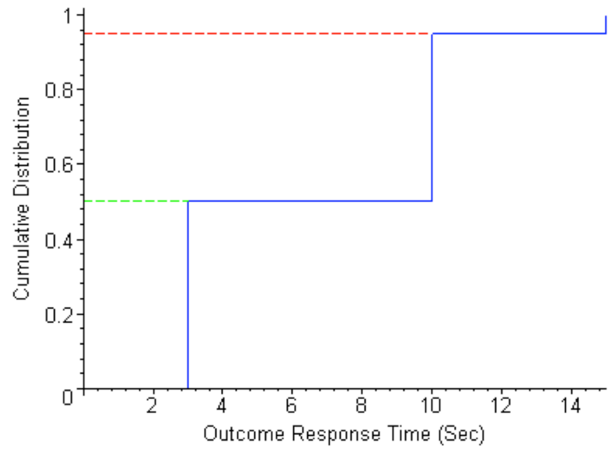

Figure 4: Defining the Desired Outcome using ΔQ

This ΔQ view gives a clear expression of what it means for a network (or other subsystem) to meet its performance requirements. For example, suppose **Figure 4** represents a desired outcome distribution:

- 50% of outcomes occur within 3s;
- 95% of outcomes occur within 10s;
- With a probability of > 99.5% the system should respond with 15s.

Figure 5 shows the measured or calculated response of the actual system, represented by the black line; since this is above and to the left of the desired outcome, the system unambiguously meets the requirement. As indicated above, this high-level view can be related to lower-level measures, and ultimately to the ΔQ of the translocation of packet flows.

Network ΔQ has several components, one of which includes delay due to signal propagation (which depends on the extent and topology of the network); another is dependence of delay on packet size (which depends on bearer technology); and a third is variability due to contention for shared resources (which depends on the pattern of demand and the scheduling of access to resources). The first two are generally constant, so the distribution of packet delays (and almost all of the packet loss) is due to the third. This is where the statistically multiplexed nature of packet networks bites, particularly when the only class of service is 'best effort' and there is consequently no basis for making more

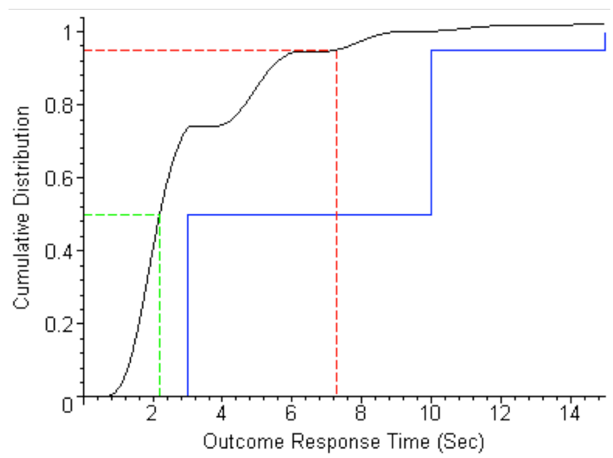

Figure 5: Satisfying an Outcome Requirement

optimal resource allocation decisions. If the ΔQ requirements of different streams are known, however, then highly optimised scheduling can be performed, delivering both predictable performance for important traffic and high efficiency. This information about the demand on the network is not usually available in the standard TCP/IP model, but in RINA it can be, due to the connection establishment protocol.

IV. Conclusion

Connecting electrical grids, urban infrastructure and everyday objects via packet networks promises many benefits. It also creates new vulnerabilities, both for citizens and for society as a whole, unless these networks are engineered much more carefully than is typically the case today. Fortunately, new scientific and technological advances are appearing that will help with this daunting task. These include the Recursive InterNetwork Architecture (RINA) and the quality attenuation calculus (ΔQ), which are approaching the necessary level of maturity due in part to work in projects such as PRISTINE. The time has come to start planning real-world pilot projects, exploiting these approaches, that can show the way to the secure, resilient, reliable and scalable 'smart networks' on which the future will depend.

References

[1] N. Stern, "Stern Review on The Economics of Climate Change (pre-publication edition). Executive Summary". HM Treasury, London, 2006.

[2] N. Komninos, "What makes cities intelligent?". In Deakin, Mark. Smart Cities: Governing, Modelling and Analysing the Transition. Taylor and Francis. p. 77, 2013. ISBN 978-1135124144.

[3] B. M. Leiner, et al., "Brief History of the Internet: The Initial Internetting Concepts" Internet Society, Retrieved 27 June 2014.

[4] J. Day, I. Matta and K. Mattar, "Networking is IPC: A guiding principle to a better Internet." In Proceedings of the 2008 ACM CoNEXT Conference. ACM, 2008

[5] B. G. Boddapati, J. Day, I. Matta, L. Chitkushev "Assessing the Security of a Clean-Slate Internet Architecture" Network Protocols (ICNP), 2012 20th IEEE International Conference on, Austin, TX, 2012, pp. 1-6. doi: 10.1109/ICNP.2012.6459947

[6] M. Ponce de Leon, "PRISTINE main achievements and exploitation prospects" http://ict-pristine.eu/wp-content/uploads/2013/12/Pristine_project_FI_factsheet.pdf 2015

[7] J. T. Bradley, "Towards reliable modelling with stochastic process algebras," PhD Thesis, Department of Computer Science, University of Bristol, Bristol BS8 1UB, UK, October, 1999.