Towards a Unified Internet Control Architecture

Philip Eardley† , Lars Eggert‡, Mat Ford†, Roger Karrer⋆, Pasi Sarolahti‡, Ralf Widera⋆, Rolf Winter⋄
†British Telecommunications plc
‡Nokia Oyj
⋆Deutsche Telekom AG
⋄NEC Europe Ltd.

***Abstract*— While the Internet has radically changed business and society over the past decades, its architecture has hardly evolved. Therefore, pressure from new applications, business models and networking technologies is distorting the architecture to a degree that requires a redesign. This paper presents two fundamental principles to redesign the Internet architecture. Firstly, the three fundamental building blocks - reachabililty, resource control and socio-economic control - must be unified to enhance flexibility while preserving its simplicity. Second, the architecture itself must be designed to evolve at run-time. The combination of these two principles will significantly enhance the reliability, robustness, manageability and functionality of the Internet, and will create new and varied business opportunities based around a common core architecture.**

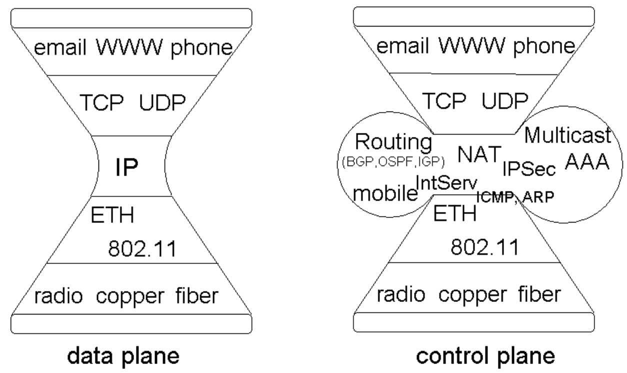

Fig. 1. Internet data and control plane (credit Jim Kurose)

I. Introduction

The Internet today is at a crossroads. Researchers, service providers, network operators and users have recognized that the current Internet architecture is ill-suited to satisfy the demands and requirements of our modern society [1]. This is no surprise: the design principles that underlie the current Internet architecture are four decades old, and the architecture has not significantly changed. Instead, new functionality and requirements have continuously been added, if not squeezed, into the architecture. As a result, the well-known IP hourglass depicted in Figure 1, with its virtues of simplicity and transparency, now only really applies to the data plane, whereas the control plane today is a waist-dominated, overloaded, ill-structured, and ossified architecture with cracks that are affecting our everyday lives.

The objective of the Trilogy project [2] is to rejuvenate the Internet architecture and make it both simple and flexible once again. Our research is based on the following two thrusts:

- unify, rather than separate, the three main components that make up the Internet architecture: reachability, resource control and socio-economic control.
- dynamically adapt to the ever increasing diversity in demand on the functionality of the Internet such as: performance diversity from Tbps core links to kbps access links; application diversity from insatiable bandwidth demand (peer-to-peer) to latency and jitter intolerant (interactive); user behaviour diversity from controlled users (company employees) to self-interested users to malicious ones, and; business model diversity from walled gardens to community networks to bundled service providers.

The research results presented herein have received support from Trilogy (http://www.trilogy-project.org), a research project (ICT-216372) partially funded by the European Community under its Seventh Framework Programme.

This paper describes these principles in detail and motivates that these two principles combine concepts that made the original Internet architecture so successful with lessons learned during the Internet's evolution.

II. A unified control architecture: the Internet hourglass for control

The cornerstones of the Internet's success have been its ubiquity and robustness: for its target applications, it provides a minimal necessary level of network support, and it self-configures (heals in response to failure) well enough to allow these applications to operate. The design principles that have enabled these features are simplicity and transparency: IP over everything, and everything over IP, resulting in the well-known hourglass shape depicted on the left in Figure 1.

However, the hourglass picture of the Internet architecture omits the mechanisms needed for control. While we have easily been able to add new protocols and integrate new technologies into the data plane, it has become increasingly difficult to add control mechanisms to the architecture. This difficulty especially holds for those requirements that were never foreseen when the design principles of the Internet were defined four decades ago - such as mobility, security or middle-box control. This is not to say that there are no control mechanisms available; but they are typically imprecise and inelegant mechanisms in the absence of suitable alternatives. Additionally, one can observe increasingly complex interactions between these individual control mechanisms. As a concrete example, consider multi-homing. Multi-homing is a simple, straightforward solution to enhance connectivity robustness and thus reachability; an implicit recognition that to build a cost-effective Internet from imperfect parts requires redundancy. Yet, multi-homing was not architected into the routing protocols, nor into the congestion control mechanisms. Therefore, neither handles is properly. To bridge this gap requires the active efforts of network operators, using crude traffic engineering tools that often serve to exacerbate routing problems.

Suitable control is becoming increasingly important and in the future it needs to be supported as a first class function, but it is far from obvious how it actually fits into the current architecture. Judging from the tremendous success of the Internet architecture's data plane, we argue that a control plane for the Internet is needed that is as simple and as flexible. It has to be an hourglass for control.

The single most critical part of the hourglass is the waist [3]. The two most fundamental services provided by the Internet are reachability (getting the packets to the destination) and resource control (sharing resources between all the sources sending their packets). As simple as this may sound, it is the basis for the wealth of applications we experience on the Internet every day ranging from e-mail to on-demand video. The lack of suitable control over both reachability and resources has made it increasingly difficult to just use this simple service. Application developers more and more often have to worry about 'the network'. Therefore, in order to regain the Internet's transparency and utility, we believe that control is needed in these two key technical areas.

A. Reachability

Reachability at first seems to be a rather trivial problem, but essentially, when looking at today's Internet and current solutions that enable reachability, it is clear that it is by no means a satisfactorily solved problem. The Internet's original requirements on reachability have significantly changed, and today's inter-domain routing is suffering from an ever increasing amount of stress: more devices need to be supported, and more functionality is added. As a result, routing protocol convergence times are in the order of minutes without even knowing which state it will actually converge to. In addition to the pure complexity of today's inter-domain routing system, requirements have changed drastically over time and, as a result, mechanisms introduced to deliver on these requirements simultaneously break reachability.

To illustrate this point a little further, middle-boxes such as network address translators (NATs) and firewalls were introduced to overcome certain problems in the reachability domain. Overcoming address scarcity and protecting a network from harmful traffic are only two of those problems. Both were essentially introduced to cope with new requirements of the evolving Internet and while the requirements changed quite significantly, the protocols did not. But what those middle-boxes do now is change the reachability assumptions of the original Internet architecture. The result is that applications built on top of it break, i.e. reachability is not necessarily

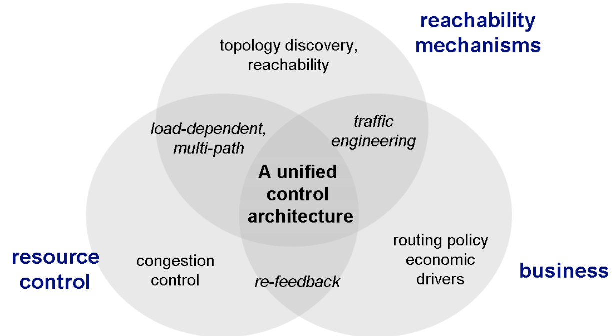

Fig. 2. A unified control architecture

achieved easily any more and certain applications need to find new and innovative ways to restore reachability which would not be their task if the Internet architecture was properly supporting the applications in the first place. Skype, for example, goes to a lot of trouble to make clients that run behind NATs reachable again and it requires some additional logic inside the network that Skype must provide.

B. Resource control

In the original, somewhat stereotyped, Internet there was a clear delineation of roles between the users who requested resources ("please transport these bits to this end point") and the providers who supplied the resources and implicitly worked out how to share resources between the users requesting them ("allocated on a best effort basis, sharing capacity according to the TCP algorithm, users all behave nicely").

However, both the users and providers have sought to expand their role by abusing the Internet's resource control architecture. Users don't just request resources but try to get more resources than they "should"; for example, by setting up multiple TCP flows, by using UDP (which does no congestion control) and by running peer-to-peer applications that have an insatiable appetite for bit transport. These are all examples of the users "tussling" into the traditional provider territory of resource control. Perhaps an even more extreme example is malicious users who launch DDoS attacks, and use these as a way of extorting money out of other parties on the Internet (another case of tussle!). On the other hand, resource providers have tussled back, and tried to re-gain control of resource allocation. For example, ISPs employ traffic policers to cap heavy users transporting more than so many megabytes in a month (or encourage these users onto a more expensive tariff). Providers also employ deep packet inspection to target restrictions on bandwidth-intensive applications, which causes users to disguise their applications, leading to an arms race between the providers and users.

III. A FLEXIBLE CONTROL ARCHITECTURE: DESIGN FOR TUSSLES

Clearly, a unified resource control architecture that would embrace the different needs of the stakeholders through a common approach would be desirable. The Internet today is a playground of many competing forces - technical, economical and social, as sketched in Figure 2. Indeed, if you ask *what is the Internet?*, the answers depend on the stakeholder: for some it is the physical infrastructure that needs to be managed, for others it is the principles and protocols that allow the Internet to work, for others again it is the enabler for services and business. The diversity in views also leads to different demands and requirements from the Internet's stakeholders.

To craft a new Internet architecture, it is necessary to understand how the architecture evolves as a function of the different stakeholders' requirements - i.e. which forces determine whether a particular solution is eventually adopted? To continue the NAT-related example of the previous section, perhaps the two main reasons for the widespread deployment of NATs aren't the technical ones mentioned earlier, but socio-economic ones. Firstly, renumbering caused by changing service provider: easy with NATs, but a real migraine with IPv6 (which is the alternative technology for tackling the shortage of IPv4 addresses). Secondly, business models: an ISP can charge business customers (who want a consistent and public IP address) more than residential customers (who are typically content to use NAT).

Widening our example to include other middle-boxes like firewalls and deep packet inspectors, a key reason for their deployment is to enable new business models like tiered bandwidth pricing and walled gardens that go beyond the primitive flat-rate subscription and settlement-free peering of the original Internet control architecture. The overall consequence is that the transparency and simplicity properties of the Internet architecture are gradually being eroded: the result is network divergence, rather than the optimistic vision of network convergence that was aspired to in the past. While there may be localised, short-term gains, at the same time the cumulative effect will be to cause the Internet to fossilise into something that supports only today's applications, and in an increasingly fragile way.

This discussion illustrates that the Internet control architecture must address not just technical requirements but also socio-economic ones. Now the basic Internet architecture is something that can only be changed very rarely - after all, the original Internet architecture has lasted over thirty years. But the architecture must be flexible and have the ability to adapt as the Internet evolves, and as the environment in which the Internet is used evolves. Therefore, we believe that central to a new Internet architecture must be the recognition that the Internet is no longer a single happy family of cooperating users and networks. Rather, the ownership, operation and use of the Internet involve contention amongst and between the various stakeholders over economic success, strategic power and social reward etc. The idea of contention (or 'tussle' to use the buzz-word) is of course familiar from the 'real world'. However [4] suggests there are some differences when considering the Internet. Chiefly these revolve around the speed at which change takes place.

In the real world of physical objects things are hard to change; for example an architect designs a new bridge many years before it's travelled over, and there are lengthy arguments (tussles) with funders, planning authorities, safety regulators and so on. However, on the Internet, capabilities are improving all the time at a rapid rate (e.g. processing power, bandwidth, network access technologies), and new applications and networking software can potentially be deployed almost instantly. Perhaps even more importantly, the value chains on the Internet are very dynamic - just think of whole new classes of very successful applications that have appeared in the last five years: social networking, VoIP, e-commerce, etc. This leads to the idea that the underlying networking protocols and architecture should allow the tussles to play out in real time: assumptions about the business model shouldn't be embedded in the design of a protocol (otherwise the protocol will get misused in order to fulfil new business requirements as they arise). Overall, by providing control mechanisms over the low level characteristics of the Internet we believe that it is possible to allow local stakeholders (the regulator, the operator, users etc.) to adjust the local settings to meet local demand, which can result in different outcomes in different jurisdictions, in different markets, for different users - to accommodate the fact that the control requirements can be different in different places and at different times.

Therefore, our key concept is to retain the ubiquity enabled by the hourglass model, and take the self-configuration philosophy one level further: **we seek a control architecture for the new Internet that can adapt in a scalable, dynamic, autonomous and robust manner to local operational and business requirements**. This self-adaptation concept caters for the need for incremental deployability and generalises the heterogeneity principles that lie in much recent architectural research.

A. An example of design for tussle - Internet resource control

We now give an example of a mechanism to update the control architecture related to resource control that utilises the properties of existing IP transport and is therefore feasible to deploy over existing networks.

Re-feedback proposes a change to the Internet feedback architecture as an attempt to design for tussle in Internet congestion control [5]. The mechanism collects congestion information along the connection path and re-inserts the information in the network along with the data traffic, allowing the network elements to know the congestion on the downstream path, i.e. between the network element and the destination. (By contrast, TCP allows the network element to know the upstream congestion.) This gives a tool which enables accountability for congestion. For example, a network operator could make a traffic source accountable for the congestion it causes when sending packets, or it could make a network accountable for the congestion it causes when forwarding packets to a peer network. But the mechanism doesn't limit the flexibility of the data sender or forwarder to apply whichever traffic patterns are required by a network application or by a network's users.

Re-feedback therefore is a good example of the design for tussle principles discussed above: the protocol mechanism does not limit the range of business models to be applied, but gives flexibility to different stakeholders in the network to apply whichever policies or behaviours best support their operation. For instance, the mechanism doesn't limit how the accountability information relates to charging. The same principles could be expanded also to other types of generic resources, which we intend to investigate in the near future.

IV. Related work

The need to design a new Internet architecture is widely recognized and reflected in national research programs such as the US FIND/GENI [6], [7] programs, of which an early summary can also be found here [8], programs in Europe within the 7th framework program, and emerging initiatives in Korea and Japan. Perhaps unsurprisingly, opinions on how to build a new Internet that is able to satisfy the demands of the next decades could not be more diverse. Here, we briefly summarise the main thrusts and highlight differences to Trilogy.

Architecture: several projects focus on architectural principles for a future Internet. Of particular interests are decomposition alternatives, ranging from mathematical modelling of algorithms and protocols over service-centric networks to architectural concepts (recursive approaches, building blocks). Moreover, novel concepts such as in-network storage and active paths are proposed.

Virtualization: virtualization enables diversity: by allowing multiple Internet architectures to be deployed in parallel, their requirements can be isolated and tailored to a specific scenario. Both approaches - virtualization and a unified architecture - have their virtues and potential, and must not necessarily be exclusive.

Economics: the lack of commercial interest in the original Internet architecture poses the question if and how economic interests can be expressed in an Internet architecture. Approaches range from exposing the coupling between economic incentives and network architectures to build a market, to market-based resource control. Trilogy shares the vision, but emphasizes the need to jointly consider economic interests with reachability and resource control.

Control and management: the overloaded waist of the control plane requires a fundamental rethinking of control and management: what are the basic building blocks to log and distil information in general and failure in particular, how can they be integrated in a simple, flexible and cost-effective manner, and who should have access to the information? Triology's main focus here lies in the joint consideration of control (reachability, resource control and economic and social control).

Routing: the routing table size and route convergence are a crucial problem in today's Internet. Approaches include updateless routing and swarming, but also questions of who decides on the path are being discussed. Trilogy is particularly interested in the integration of multi-path and multi-homing into a unified architecture.

Security, Privacy, Identity: the integration of personal devices with ubiquitous access poses significant problems. Solutions such as virtual identifiers, the separation between location and identification, the definition of private name spaces and using attributes to protect privacy provide novel approaches towards a higher protection of users.

Others: This category contains proposals for novel data plane technologies, such as wireless, sensor and optical networks. They are orthogonal to Trilogy since Trilogy focuses on the control plane.

V. Conclusions

Independent of the proposed design and enhancements to the architecture, and independent of whether these projects start with a clean sheet or recognize today's Internet as a starting point, the re-design of the Internet is vital to the Internet's success, both in terms of its fundamental principles as well as its subsequent evolution. Only by understanding these principles, with their failures, shortcomings and successes, can we craft a new Internet architecture.

We have identified two core principles for a future Internet architecture: a unified architecture and an architecture that dynamically adapts to current and upcoming tussles. Both principles together will lead to a future Internet architecture that combines simplicity with the much-needed flexibility in the presence of future changes.

References

[1] M. Handley, "Why the internet only just works," *BT Technology Journal*, vol. 24, no. 3, Jul. 2006.
[2] "Trilogy project," http://www.trilogy-project.org/.
[3] J. Kurose, "Networking success: new challenges, and an expanding waist as the field approaches 40," IEEE Infocom 2004 Keynote talk.
[4] D. Clark, J. Wroclawski, K. Sollins, and R. Braden, "Tussle in cyberspace: Defining tomorrow's internet," *IEEE/ACM Transactions on Networking*, vol. 13, no. 3, pp. 462–475, Jun. 2005.
[5] B. Briscoe, A. Jacquet, C. D. Cairano, A. Savatori, A. Soppera, and M. Koyabe, "Policing congestion response in an internetwork using re-feedback," in *Proc. ACM Sigcomm*, Aug. 2005.
[6] "Global Environment for Network Innovation (GENI)," http://www.geni.net/.
[7] "Future Internet Design (FIND)," http://www.nets-find.net/.
[8] A. Feldmann, "Internet clean-slate design: What and why," *Sigcomm CCR online*, Jun. 2007.