Zen and the Art of Network Research

Jon Crowcroft, Tolga Uzuner and Andrew Warfield
University of Cambridge Computer Laboratory
15 JJ Thomson Avenue, Cambridge CB3 0FD, UK
{firstname.lastname}@cl.cam.ac.uk

July 15, 2003

"In the sound of the bell of the Gion Temple echoes the impermanence of all things. The pale hue of the flowers of the teak-tree show the truth that those who prosper must fall. The proud ones do not last long, but vanish like a spring night's dream. And the mighty ones too will perish in the end, like dust before the wind."

– From the opening lines of the "Tale of the Heike"

1 Introduction

It seems to be increasingly clear that, in regard to the design of global data networks, an era of networking is coming to an end. Suddenly, what was once a playground for researchers now supports both culture and commerce, and has increasingly become a part of our daily lives. The present-day Internet quite clearly serves a basic infrastructure role, alongside other communication and transportation systems. As such, its ongoing evolution will inevitably be the result of a tug-of-war between the vested interests of not only academics, but corporations, governments, and consumers. Gone are the days of sudden, drastic change.

As academic network researchers, it seems sensible to consider the future directions of our field. After some thought, we feel that several points are worth raising.

Although the Internet has changed drastically over the past ten years, the mindset of researchers perhaps has not. The so-called ossification of the network is the inevitable result of a large body of applications that depend on interfaces remaining static. Moreover, as the success of the network continues, this cloud of legacy can only continue to inhibit anything but incremental change.

The evolution of emerging markets have been studied, seminally by Cohen [4], who observed that new markets are typically characterised by "product innovation", in which a new product emerges, and is later followed by "process innovation", where incremental improvements are made on an established market. The nature of these later innovations is one of optimisation and streamlining, more commonly the domain of corporate rather than academic research interest. The Internet, we feel, has reached this point of maturity.

This line of argument is not intended to declare the end of academic networking research, but rather to help advise the direction that such research should take. We suggest that the incremental improvement of the existing artifact be left largely to the corporations who now maintain it, and that academic research focus on the science of the existing network, and the architecture of the networks of the future.

The remainder of this paper argues first that network research now has a broad realm of complementary research to draw on, and that as researchers we must endeavour to become interdisciplinarian. We go on to propose that the future of research for global networks has two fundamentally different directions: First is research concerned with formalising our understanding of the present Internet, and second is the task of developing the network that will follow, an Internet to last for another hundred years.

2 Science and Innovation

The Internet has a history of learning through experience. Throughout its growth, there simply were no comparable systems that could be drawn on for underlying models. The approaches taken to control the network and deliver data have been established through experimentation, evolution, and cautious deployment.

Although these rudimentary approaches have served the development of the network well thus far, many neighbouring fields have matured to the point that they have plenty to offer in terms of composing an underlying science upon which the network can evolve.

For instance, our understanding of nonlinear control, and the complexity of large-scale systems are finally maturing. As an example, in the last decade, while it has been shown that many control problems are NP-hard, "good enough" solutions are becoming available [1].

In the area of complexity, much has been done on emergent complexity but not a lot of it has been applied systematically to the area of networking: while many issues remain unresolved and it would be hard to say that an overall theory of complexity exists, one can at least say that a lot has been done to define what complexity is [6] and how to measure it. Specifically, in one recent count, we found over 70 different "notions" or ways to measure complexity, be it algorithmic, information theoretic, or effective. It is time to apply these techniques to the networking of the 90's to figure out the way in which these networks are "complex" and also how this kind of complexity matters (i.e. cost and difficulty to operate).

These are hardly the only related fields. There is increasingly an established corpus of techniques to understand and interact with large, complex systems. Theoretical physics, machine learning, and economics all have additional contributions to make in this domain. Of these, we focus on two specific areas from which we expect these insights to come: manufacturing and finance. Our reason for looking at these fields for insight is straightforward: they have been under the discipline of market competition for centuries. Telecommunications on the other hand has been a truly competitive market for less than a decade.

A similar analysis to what we propose was applied to manufacturing systems between the late seventies and the early nineties, of which Casti's and Golay's studies are among the most seminal [2, 7]. By applying a similar, thorough analysis of networking, practitioners will be able to talk systematically and meaningfully about the systems they architect.

We see scope for information transfer from everything that is known in the manufacturing world to telecommunications in at least in eight different areas:

1. Flows: Manufacturing experts recognise the importance of flows, with respect to events and outcomes, their trajectories over time and the relevant production processes, which drive those events.
2. Strategy: The recognition of complexity as a strategic core advantage. The roles of mix, indices and metrics for structural, operational, and static complexity.
3. Product Strategies: Product strategies, rationalising the product mix, product families, bundling strategies. Component strategies.
4. Facility Design: Assembly lines, plant configuration, calculating capacity and technology requirements.

5. Perational complexity: The analysis of existing complexity: programmable and non-programmable states, queue behaviour, operational complexity. Linking structure to operational performance: Why queues form, predictable and stochastic queues in operations.
6. Control complexity: Levels of tolerance, and trajectories. Impact of Disturbances, preventative measures, sources, handling, minimising impact, and recovering.
7. Planning: The role of planning, scheduling, materials requirement planning, synchronised manufacturing.
8. Supply chain expertise and technologies: Parallels and differences with factory systems. Simplifying and controlling supply chains. Make or buy decision rationale. The roles and principles of partnering. Exporting complexity through the supply chain.

From the field of finance, we also believe that there are exciting opportunities to transfer know-how into the field of telecommunications, and we are not talking about simply the financial management of a service provider. We believe the applications can be much more low-level and analytical.

We see three opportunities:

1. Modelling: Asset returns in financial markets and traffic flows are very similar: power laws, fat-tails, highly skewed distributions, and stochastic volatility. It would be fair to say that the quantitative analysis of distributions such as these is most advanced in the markets for derivatives, for obvious incentive reasons. Of these, the credit derivatives markets have gone furthest in modelling correlations amongst large numbers of variables and in modelling the associated distributions.
2. Option Theory and "Real Options": Since 1973, the financial markets have been trading a concept called volatility, thanks to Fischer Black, Myron Scholes and Merton Miller. Most of the decisions made by telecommunications systems and operators are the same as those made by option traders: they are made under uncertainty, when there are many alternatives to a single solution, decisions lead to further indecisions, and outcomes are contingent on other variables.
3. Portfolio Theory: How does a portfolio manager decide the right mix of assets, from amongst millions, and under constraints? This know-how is directly applicable in the area of integrated management of a large-scale telecommunication system, whether its the right mix of flows, equipment, or the right configuration.

Finally, as a science is developed to describe the theoretical underpinning of the Internet, we must also endeavour to understand and provide techniques to support the overlying administration of such networks. The management of these networks has largely been ignored within existing research, while the interactions of large organisations with frequently conflicting interests is a fundamental issue in the network today.

3 Chautauqua

The construction of a solid theoretical underpinning to understand the Internet is vital in order to allow further incremental improvements to the current architecture. Not only must a set of applicable theory be composed, as mentioned above, by interacting with other disciplines, but these results must be incorporated into the existing set of tools for modelling the network.

Through the development of such an underlying science, and the incorporation into the existing tool-set, we hope that researchers will be more able to achieve reproducible results. The fundamental difficulty here is in the effective description of the complex system state that invariably underlies network experiments. As such, theoretic work will need to provide simulators and emulators with mechanisms to characterise complex environment state. Moreover, tools will be required to accurately capture the current state of the network, in order that researchers may reproduce and share specific observed behaviours.

In this regard, networking research seems somewhat like other physical sciences: as we progress, we become aware of the need for more fine-grained and fundamental mechanisms of both observation and inference.

4 Motorcycles

We now have a motorcycle, it can go pretty fast but we also know it can't fly. Before moving on to design an airplane, let's make sure we leave the manufacturers with a good model of how it works and what it's current failures are. And move on.

The shortcomings of IP have been discussed before [8], by filling in a body of theory to underly the existing system, it's maintainers should be well able to carry it forward.

We propose that the next great academic problem in networking is the design of what comes next. It seems clear to most researchers that it is unlikely that the existing Internet, even with an ongoing set of feature additions and upgrades, is the eternal data network. Fortunately, unlike many other large infrastructures such as railway or postal services, the structure of the Internet *can* be changed quite drastically at a low relative cost. Routers and software may be upgraded or replaced while still using the same underlying physical media.

Much research effort to evolve the existing network is presently being invested in the creation of services atop overlay networks [9]. While this is a very interesting area, and considerable novel research has emerged, it is highly unlikely that architectures to rival the existing Internet will emerge in an environment where they must use the Internet as an underlying layer.

If it is in fact in the interest of the research community to construct so-called disruptive technologies, they must be enabled, as they were at the onset of the Internet itself, with a low-level network testbed, upon which they may explore approaches to network architecture without being inhibited by the existing architecture.

In addition to the need for such an infrastructure, and for research that does not take the existing Internet as a precondition for its validity, there is a final area of research which is crucial, but largely unexplored at present. Assuming that there will eventually be some non-incremental change to the existing network, there is a need to understand and provide mechanisms to enable a transition from the existing network structure to whatever comes next. The lack of research in this arena is clearly evidenced by the inability to successfully deploy new protocols, such as IPv6 within the existing network. Providing generalised mechanisms to facilitate the deployment of new research should be a key goal in the near future.

For these important reasons, we argue that future networking research should not focus on the Internet! We feel that we can make this proposition because, we believe that an architecture exists whereby heterogeneous networks can co-exist, making it possible to carry out research on non-Internet networks which have a fair chance of "making it in the real world".

In "Plutarch: An Argument for Network Pluralism" [5], Crowcroft et al propose a new internetworking architecture which subsumes existing architectures, such as the Internet Protocol. The proposal is built around two important notions: contexts, defined by environments in which addressing, naming, routing and transport were common, and intersititial functions, which allow messages to transit across contexts.

The key significance of this proposal in the discussion on the future of networking research is that it makes it possible and believable that research in new networks and around new mechanisms for addressing, naming, routing, and transport can be implemented in the real world in the future, even, and especially, if this does not happen in the Internet as we know it today.

New network architectures can co-exist with the Internet, and perhaps even bypass or supercede the Internet in the future. This is done by embracing heterogeneity in the hope of allowing radical innovation. The homogeneous Internet architecture and its advantages are not abandoned but retained as one architecture among many.

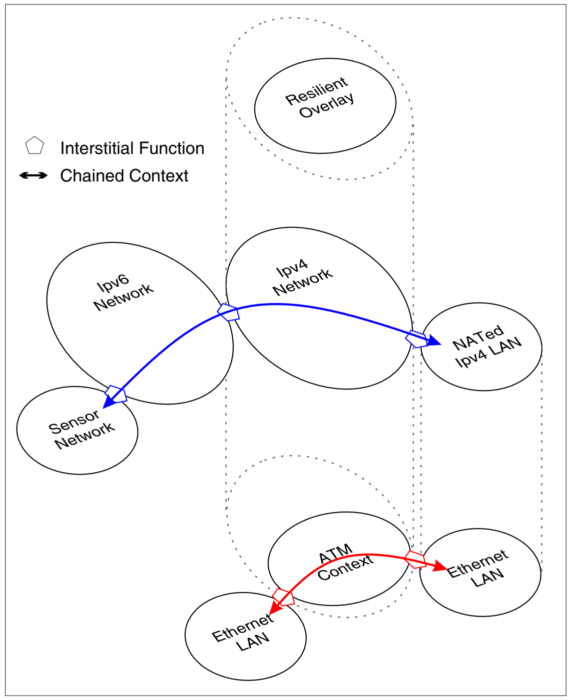

Figure 1: Communicating across hetrogenous networks in Plutarch

Frankly, this viewpoint is in direct contrast to those presented recently at HotNets-02, centered around the use of overlays as the basis for networking research. Such research is bound to be limited by the many limitations of the Internet today, all too numerous to count. Overlays are not a significant enough step in the search algorithm to be able to move out of the local minima we know as the Internet today: eventually, the deadweight of the existing infrastructure will drag one back to the lowest common denominator known today as IPv4, BGP and DNS. We believe that scientifically significant, and disruptive research can only happen in new contexts.

We first proceed to explain very briefly how this new architecture can work. Most of this description is taken verbatim from the Plutarch paper mentioned above.

We then go on to explain our motivations for this new architecture, again primarily from the Plutarch article, but further argue that in the future, the notions introduced by Plutarch have the potential to significantly disrupt the Internet, by introducing new contexts that wholly replace the Internet by triggering whole-scale migrations.

4.1 Architecture

In Plutarch we divide the world into contexts, each comprising some set of hosts, routers, switches, network links and so forth. Within a context we expect homogeneity regarding such things as addresses, packet formats, transport protocols and naming services. Distinct contexts differ in at least one of these areas.

Communication across a set of contexts is enabled by interstitial functions, which map between the sets of functionalities encapsulated by contexts. We can divide such functionalities into four main areas:

- Addressing: with exposed programmatic interfaces.
- Naming: alternative approaches of mapping between a plurality of naming systems as in VoIP, for reasons of scalability and administrative overhead, as opposed to just DNS
- Routing: different styles of routing protocol as appropriate in different networks.
- Transport: optimized transport protocols for specific network types

Functions such as naming, addressing, routing and transport must be supported end-to- end across radically heterogeneous networks through the addition of suitable explicit interaction at boundaries.

By making these regime transitions explicit, we believe that (i) the network model will more accurately reflect the network's reality, and (ii) the network model will be more extensible, allowing new services to more easily be incorporated at all architectural layers.

4.1.1 End-to-End naming and addressing

In our view, the twin functions of naming and addressing should be implemented in accordance with the end-to-end argument. Paradoxically, the current Internet imposes single mechanisms for addressing (IPv4, by design) and naming (DNS, by an accident of evolution) from the middle of the network. This leads to a requirement for globally-bound names and addresses. The model is already insufficient to capture the architecture of the current Internet (NATs, IPv4/v6 gateways, dynamic DNS servers, etc.), and does not address the connection of other, radically different networks (sensor nets, planetary-scale overlays, etc.).

The central problems in this scheme are communicating and resolving names and addresses across network boundaries. The goal of our architecture is to provide a set of compositional building blocks that may be used to allow the composition of heterogeneous networks in order to provide an end-to-end service. The two abstractions we propose are the context and the interstitial function (IF).

4.1.2 Contexts

Contexts serve two purposes: first, they describe communication mechanisms embodied by different networks, within which an endpoint might bind a particular communication session; second, they serve as descriptors allowing end-to-end services to be composed via the application of network closures.

4.1.3 Interstitial Functions

In order to accommodate the differences between contexts, while still providing an end-to-end service, data may have to be manipulated at context borders. To achieve this goal, we introduce the Interstitial Function (IF), whose purpose is to allow data to pass between two adjoining contexts. Contemporary examples of IFs include NAT boxes, signaling gateways, and BGP routers. However, we also envisage situations where IFs may explicitly be used to bridge dissimilar transport networks (e.g. IPv4 onto ATM) or to provide high-level service modification, such as transcoding video streams or inserting forward error correction on unreliable links.

4.2 Motivation

4.2.1 More Competition

In essence, Plutarch allows different contexts to compete with each other for business, but on a more level playing field then is currently possible in the context of the Internet. For example, it is simply not possible to change the naming schema in the Internet today within the Internet, even if a superior one were available. One cannot fight the Matrix purely from within!

However, if a suitable architecture for passing traffic from one context to another is made possible, then one could envisage building new contexts which significantly outperform the Internet on the basis of one or more of the four variables mentioned above: addressing, naming, routing, and transport. Over time, nodes which benefited specifically from the performance boost of one context would migrate over. At the same time, with the use of Interstitial Functions, it would be possible to continue servicing customers in the old Internet. As customers found interaction in the new context significantly outperformed the old Internet, they too would switch over to the new context, as the Interstitial functions would surely introduce some form of delay. Furthermore, one of the key benefits of the new context may be reduced complexity in administration or management as it relates to specific transactions, itself a reason to switch contexts even if there were no traffic penalties.

There are a large number of reasons why wireless networks would benefit significantly from a new context, and one could see a significant migration from todays Internet to one optimized around on-demand protocols. To take an example from [5], connecting an ad-hoc wireless network with on-demand routing to an Internet AS running OSPF and BGP requires a more complex mapping than simply exposing the routes of each network via BGP. Using an explicit interstitial function to map routing information between these networks hides the churn of Internet routes from the on-demand protocol, deals with the wireless network protocol's broadcast, and hides the instability and mutability of routes within the wireless network.

4.2.2 Architectural Reasons

We also believe that the Plutarch framework makes sense for a number of architectural reasons. The principal one is the concrete problem of connecting networks where a common overlay protocol such as IPv4 or IPv6 is infeasible or undesirable, for example sensor networks, or specialized networks which offer valuable intra-domain functionality which IP must ignore. We also believe that Plutarch captures the state of the Internet we see today better than models based on the Internet's original architectural principles [3]. Finally, a model of networking based on explicit contexts provides a clearer framework within which to debate future architectural changes to the Internet than the current tradition allows. We hope it will allow the debate to move forward.

5 Conclusion

As the Internet matures and increasingly acts as a supportive technology for modern society, the field of research that created it is presently coming of age. As academic researchers, our research testbed is now fraught with corporate and cultural inertia.

In this paper, we have explained our belief that there are now two fundamental directions in which academic networking research should progress. With regard to the current network, we pose that researchers be less concerned with incremental changes to the current architecture, and more with developing a thorough science to underpin and explain its operation. To achieve this we advise an interdisciplinary approach that applies the present state of knowledge in many surrounding fields.

The second direction in which we feel research should advance is in the architecture of the next global data network. As the Internet cannot possibly be extended to provide for the scale or the breadth of requirements that it will face in the distant future, researchers must endeavour to consider what a network that can face such challenges will look like. We argue that in order to facilitate the exploration of such potentially disruptive research, the community will need a low-level testbed, similar to what the early Internet was, in order to experiment with new things. Additionally, we suggest that in order to facilitate the adoption of the results of future research, we will need to develop mechanisms to allow new technologies to be more easily deployed within the network.

The future of networking research is, in our opinion, a bright one. In the years to come, we anticipate both a maturation of fundamental principles, as networking comes into its own as a science, and also a wealth of new innovation as we look to the future.

References

[1] BLONDEL, V. D., AND TSITSIKLIS, J. N. A Survey of Computational Complexity Results in Systems and Control, 2000.

[2] CASTI, J. On System Complexity: Identification, Measurement and Management. In *Complexity Language and Life: Mathematical Approaches*, J. Casti and A. Karlquist, Eds. Springer, Berlin, 1986, pp. 146–173.

[3] CLARK, D. The Design Philosophy of the DARPA Internet Protocols. *Computer Communication Review 18*, 4 (Aug. 1988), 106–114. Proceedings of ACM SIGCOMM 1988.

[4] COHEN, W. M., AND LEVIN, R. Empirical Studies of Innovation and Market Structure. In *Handbook of industrial organization* (1989), R. Schmalensee and R. D. Willig, Eds., vol. 2, pp. 1059–1107.

[5] CROWCROFT, J., HAND, S., MORTIER, R., ROSCOE, T., AND WARFIELD, A. Plutarch: An Argument for Network Pluralism. *To Appear, SIGCOMM Workshop on Future Directions in Network Architecture (FDNA)* (August 2003).

[6] GELL-MANN, M. What is Complexity? *Complexity 1* (1995), 16–19.

[7] GOLAY, M., SEONG, P., AND MANNO, V. A Measure of the Difficulty of System Diagnosis and its Relationship to Complexity. *International Journal of General Systems 16* (1989), 1–23.

[8] MOLINERO-FERNÁNDEZ, P., MCKEOWN, N., AND ZHANG, H. Is IP Going to Take Over the World (of Communications)? In *ACM HotNets-I* (Princeton, NJ, October 2002).

[9] PETERSON, L., CULLER, D., AND ANDERSON, T. PlanetLab: A Testbed for Developing and Deploying Network Services, June 2002. Whitepaper, available at `http://www.planet-lab.org/pubs/vision.pdf`.