FIE - Future Internet Enervation

Jon Crowcroft
The Computer Laboratory, University of Cambridge
Cambridge, UK
jon.crowcroft@cl.cam.ac.uk

This article is an editorial note submitted to CCR. It has NOT been peer reviewed. The author takes full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

I'm so Bored of the Future Internet (FI). There are so many initiatives to look at the Internet's Future[1], anyone would think that there was some tremendous threat like global warming, about to bring about its immediate demise, and that this would bring civilisation crashing down around our ears.

The Internet has a great future behind it, of course. However, my thesis is that the Future Internet is about as relevant as Anthropogenic Global Warming (AGW), in the way it is being used to support various inappropriate activities. Remember that the start of all this was not the exhaustion of IPv4 address space, or the incredibly slow convergence time of BGP routes, or the problem of scaling router memory for FIBs. It was the US research community reacting to a minor (as in parochial) temporary problem of funding in Communications due to slow down within NSF and differing agendas within DARPA.

It is not necessary to invoke all the hype and hysteria - it is both necessary and sufficient to talk about sustainable energy[2], and good technical communications research, development, deployment and operations.

To continue the analogy between FI and AGW, what we really do not need is yet more climatologists with dodgy data curation methodologies (or ethnographers studying Internet governance).

What we do need is some solid engineering, to address a number of problems the Internet has. However, this is in fact happening, and would not stop happening if the entire Future Internet flagship was kidnapped by aliens.

"We don't need no" government agency doing top down dictats about what to do when. It won't work and it will be a massive waste of time, energy and other resources - i.e. like AGW, it will be a load of hot air:)

On the other hand, there *are* a number of deeper lessons from the Internet Architecture which might prove useful in other domains, and in the bulk of this opinion piece, I give examples of these, applying the Postel and End-to-end principles to transport, energy, government information/vices.

[1] http://www.future-internet.eu/, http://www.nets-find.net/ and similar programs in pretty every any other geo-political arena

[2] See for example David Mackay's Without Hot Air book, at http://www.withouthotair.com/

Categories and Subject Descriptors

C [.]: 2.1 [Packet-switching networks],[Network Communication]

General Terms

General Terms: Algorithms, Performance, Design

Keywords

Communications Systems Research. The Internet

1. INTRODUCTION

I believe that we can consider that the Future of Internet Network Research is in hand.

There are a number of important challenges for the Internet today, but they are things that are being addressed and delivered by ordinary commercial pressure on competing ISPs, on equipment and software vendors, and on/from application service providers (aka Cloud). Where challenges are beyond the horizon for these organisations, there is sufficient innate curiosity in the academic and industry research labs to tackle the problems. Indeed, I would claim that a strong personal motive for some academics is to *have a visible impact*, and that there is no finer and better way to do so than to be relevant to industry[3]. The Internet is so visible, and its problems so blatant that a moments reading and studying reveal a plethora of fascinating things to tackle at the level of undergraduate projects, Masters and PhD students' work, and programmes of work for whole labs. No special initiatives from top down are needed at all.

Of course, we are familiar with the notion that the Internet, like food, transport and energy, is a critical resource. It is, as Neal Stephenson uses the word, *Infrastructure*.

Hence, it is required to exhibit some properties we associate with 21st century infrastructures. We would like the Internet to have resilience in the face of component failures, to continue scale in cost (measured in money, or in complexity of memory, processing, storage, etc) sub-linearly in the number of nodes/users/objects, and to retain its immense

[3] Note that I do not subscribe to impact[4] as the sole and ruling metric for whether work is important - foundational work is impossible to assess on the basis of impact; and the pure curiosity-based (or other more mysterious) motives of people to work on foundational problems are quite other than those of people wishing to be the "father of the Interweb", and should be equally, if not more highly applauded.

flexibility as the demands made grow in scope and diversity of applications and performance.

We know that we must confront the continuing excesses of end users untrammelled desire for content distribution, whether as user contributed publication, or simply demand for all the output of human creativity at any given subscriber. We all see the emergence of the Internet of things, where sensors and actuators are connected to the web, creating another leap in demand for addresses and another step in the requirement for security, privacy and resource control. We can all see that the transmission and switching technologies are heading towards all optical core and fixed access networks, with wireless access already overtaking the fixed access in sheer numbers of users.

Etc etc etc; Yada yada yada Blah blah blah.

All of these facts have been the case since the Internet was first called that. Plenty of cool research is being done to address these problems - I'd cite projects like Trilogy working on resource pooling and multipath/multi-homing; projects abound working on low cost optics, on next generation wireless, on sensor nets; on improved privacy; on more transparent management of cloud services; and so on.

Evolution rather than Revolution[5] will take care of things, provided we don't stall the funding of good communications research (again).

What might be an alternative playing field for those that want to engage in these top-down initiatives?

Here is my idea[6].

Why not use the Internet as a model for other public services?

The Past Internet delivered some powerful and persuasive lessons in:

Decentralisation A thoroughly decentralised system has a number of properties in terms of scaling and resilience. Decentralisation allows local management independent of global, which means diversity and heterogeneity are naturally supported.

Federation The Internet is highly federated. There are five orders of magnitude more service providers at the network than there were in 1992 when it was first divested from US government funding. At other levels (e.g. application/datacenter and cloud) it is even more rich, and yet not complicated.

Adaptation The Internet adapts. Protocols adapt to availability of links, paths, routes, interfaces, capacity, and so on. Maybe this is "just" good engineering (many feedback loops) or perhaps it is something more (Highly Optimised Tolerance?). It is certainly effective.

Evolution The Internet evolves. From the seventies til the late eighties, it supported e-mail and file transfer and remote terminal access. From the early nineties, file sharing and the web. In the noughties, video streaming, voice, social networks, cloud computing and more. And yet an IP packet from 1980 could still make its way from source to destination quite happily, unscathed.

[5] See the excellent recent paper by Dovrolis at `http://www.cc.gatech.edu/fac/Constantinos.Dovrolis/Papers/ccr-evolvable.pdf`

[6] I'm not the only person saying this and so I am not claiming to be the originator of this at all

How was this be achieved? Past Internet lessons revolve around certain characteristics (this list is not exhaustive, and others have elsewhere documented both characteristics and principles, better than I):

Decentralised, Hierarchical Routing The Internet routes around trouble. Routing is also opportunistic (some people say greedy), and takes advantage of new paths as they arise. However, the hierarchy in the net means that these tropes occur without global coordination being necessary.

The Internet pools resources. Packet switching permits statistical multiplexing. This is a hugely more general concept than the simple avoidance of fixed resource allocation and isolation that underpins circuit switching. Resource pooling is a form of *mutuality.* If I don't use a resource, why should I stop you from using it? I shouldn't, if that means that when you don't use a resource, I can use it. This mutuality needs some light, minimal enforcement (to avoid trivial commons-mode failures), but the enforcement is surprising - it rests in the next item on this list.

Distributed Resource Management Congestion control in end systems is a decentralised form of resource management, and takes advantage of end users and applications intelligence, (a.k.a. "showing some adaptation").

The meta-rule is that you *must* implement some form of adaptation, or congestion control, in the face of feedback. This is the only rule necessary, when coupled with the dynamic routing (and hopefully in future with multipath routing) to allow proper and sufficient pooling of resources.

Federation The Border Gateway system implements a rich policy routing system that admits of many business models. By being un-restrictive about these business models, it allows evolution of new forms of service provision, together with interworking between diverse forms of service provision. It also allows resource pooling, but with some minor constraints (service providers can choose whether they are a local or more global player in this game).

Layered Abstraction Almost prior to the Internet, the idea of abstraction, and in particular layered abstraction had taken hold in Computer Science. It goes hand in hand with a network, horizontal relationships implemented by open *protocols*, and a modularised system, vertical relationships between modules implemented by open *interfaces*, to provide an almost infinitely flexible system.

Layers and modules can evolve independently.

Hence the Web, VOIP, IPTV, OSNs, etc

Advertisement/discovery Adding a component, whether it is a service or a host, or a router or a network is trivial. The system is based on "soft state" - we do not have central registries.[7]

[7] Some crucial identifiers do need a central agency to prevent collisions of allocation, however, that should, ideally, be the sole role of a registry. Hierarchy and tagging usually suffice

Rendezvous No pre-arrangement necessary - this could also be seen as *any source unicast*. This is a boon and a blight - obviously unsolicited content is a pain. However there are scaleable Internet approaches to tackling this.

Hierarchy Systems Scale arbitrarily through hierarchy. We use this for names, addresses, routes and capacity. However, there are interesting problems with hierarchies that force the designer to make a choice - e.g. for identifiers and routes, we must pick usually one of the list: topological, organisational or geographic. That choice has consequences for users that want scaling in a different axis. This is one of the continuing areas of confusion and difficulty in Internet architectural research, showing up in many places (not least, multicast, mobile, and accountable Internet addressing support).

The Internet is not just a set of engineering artefacts. The intuitions of its creators (who are many and often unsung) turned out to contain Organisational Principles. These have been teased out, or elicited by observation and through many examples of their applications.

Architectural Principles include the following (incomplete list):

Abstraction Computational Thinking provides us with the powerful idea of abstraction. Hiding details in design and in implementation is a neat trick. When applied to a network of hierarchically composed systems, it leads to a set of designs which exhibit the desirable properties we outlined above.

End-to-end model The concept of De Minimis/Parsimony is at least two thousand years old. Its application in the Internet is rather more recent. In a layered, networked system, the principle leads to choices about details of functionality in the architecture (where to put resource pooling, how to control routers, addresses, congestion, etc). The principle doesn't lead to a static design, but leads to a flexible architecture which can have an evolutionary implementation.

Postel Principle Robustness is exhibited in a lot of places in the net. This is because a guiding principle is to "plan for the worst, but hope for the best" when designing and implementing any component of any protocol. This is also good computer science (handling exceptions:).

Amdahl/Cray/Moore/Metcalf/Reed's laws etc There are a lot of observations about the growth scale of the net over time and the value of the net at a given time. Most are super-linear, which allows a lot of slack in the design.

Consumer electronics scale The slack in the design matters because as we go through each step of evolution of the net, we need to scale up the production of components (e.g. low cost home routers when ADSL arrived, or fancier batteries for smart phones , or low cost WiFi chips). As with big Pharma, the investment in production technologies for these technologies is pretty big, so you need a pretty big, big bang sales to make it pay off. One next step, low cost photonics for fibre-to-the-home needs a clear story of the right kind, right now.

to allow local use of the same name for different contexts. viz `apple.co.uk` and `apple.com`

2. THE FUTURE INTERNET AS A MODEL FOR FUTURE PUBLIC SYSTEMS

So instead of worrying about the future of the Internet, lets worry about the future of everything.

Crucial public systems deliver food, energy, transport, housing and so on. Three systems seem like low-hanging fruit when it comes to re-application of the ideas behind Internet:

Transportation In transportation, we largely plan a journey and all its component path segments before the trip. This is analogous to the circuit switching model. Indeed, it is very crucial to safe operating of scare resources like fuel and landing slots for air travel, or for careful segregation of trains when close to points ("switches", un-ironically, in American English).

However, modern control systems might mean that we could consider a more dynamic approach to transport, perhaps using the end2end principle. Resource pooling (car ride sharing for example) is a clear win here too.

Energy Energy has traditionally been generated by a small number of providers, and delivered by an even smaller number of national grids. However, the model of local, sustainable production of (or return of unused) energy means that we need to consider many more business models and many more interconnect patterns - perhaps there are lessons here from the Internet, and we might consider using its federation model of providers. Routing around trouble and resource pooling are clear potential wins here too.

Government Governments are big and centralised. This is an inefficient way to organise what is essentially an information brokering service. Perhaps the Internet and the web offer an extreme decentralisation model, with abstract interfaces and protocols that would scale government to the small, personal level and to the large, far more effectively.

2.1 Transport

I would advocate the use of the end2end principle in the design of future transport systems.

Firstly, we need to separate out routing and `transport`. Treat the person (or parcel) as a packet, switched to destination through a sequence of steps, not necessarily predetermined.

Then the user or provider are free to adapt routes *in situ*, or, so to speak *en route*, taking advantage of resource pooling.

The use of end-to-end resource management mechanisms such as congestion control are strongly indicated. The notion of *congestion exposure*, through publication of current resource usage, is clearly the way forward in the Internet and on the roads and in the air. [8]

[8] The London congestion charge was therefore a good start.

2.2 Energy

We envisage potentially as many energy providers as there are homes, people or vehicles. Electric cars storing charge might actually be buffers for energy, uploading it to neighbours depending on price. Turbines or other ways of generating energy in a renewable manner often provide power when the demand is not there. However, by moving resource demands around more dynamically, we can take advantage of this. For example, data centers can be distributed and replicated wherever there are intermittent energy sources, and bought on stream whenever supply of electricity exceeds local demand. We can trade latency (for users of information) for efficiency (in use of intermittent supply).

This requires efficiency and flexible means to federate a huge number of providers. For this to work, we cannot assume any particular set of business models. Instead, like the Inter-domain routing system of the Internet, we should build a subtle meta-business rule system.

As with traffic, we can also take advantage of other Internet ideas such as adaptation (a la routing and traffic engineering). We can use higher level information about demand, and we co-optimise Energy with other things E.g. 1. Transport & 2. Interweb

2.3 Government

The UK government is as centralised as you can get. The use of information systems by government has sadly led to more centralisation, and yet not necessarily to any improvement in government services.

The creeping *database state* is a Google-style paradigm [9], where the social benefits of data-mining of all personal information is envisaged as outweighing the disadvantages of all the loss of personal privacy and all the risks that that entails.

However, this is entirely unnecessary to achieve the social goals that are sought. My bank does not need to know my health record. My car insurance company does not need to know my voting preferences. My social network provider does not need to know my physical postal address. There are many many providers and only I personally should be the key holder for each. Not only is a separation of concerns desirable for privacy reasons, it achieves lower cost, scaleable and evolvable systems. It is also the case that privacy preservation technologies can be deployed with greater success on a separate set of systems than on one monolithic basket-case[10].

The Internet principle of decentralisation is equivalent to the principle of subsidiarity, once held up as a founding principle for the EU (and implicit in the name of the Federation that is the USA). IP name/address/content governance all contain this principle too. Hierarchy and abstraction hides complexity, and does not require centralisation at all. An ideal Internet-style government would be nothing, but a set of common protocols and interfaces.

We have said that hierarchies may be inflexible in the face of multi-dimensional requirements (geography, topology, provider). This needs more subtle thought. I do not believe it is insurmountable.

[9] The latest instantiation is the *gCloud*, whose early steps can be seen in the `data.gov.uk` service.

[10] Nothing is perfect, but certainly the risk reduction is worth it.

3. ARCHITECTURE RENEWAL

A colleague pointed out recently that IEN-1, written at UCL in 1977 before any RFCs were started, says "Internet routing would appear to be far more difficult than is currently believed". This is in the context of trying to figure out what the future addressing system should look like. IP addresses are overloaded[11] in many marvellous and creative ways - they are truly polymorphically perverse.

Just taking the use of addresses as an example, and looking at the way that functionality has accreted in the Internet around, one can see a process that is very common in large systems. There is mission creep, there is increasing entropy, there are complications due to the steady (linear) increase in shims leading to an emergency appearance of an overall rickety, Heath-Robinson machine of alarming fragility. However, this appearance in the eye of those adherents to the original architecture, does not necessarily indicate that the system needs a re-design.

Any system that lasts this long will go through this process. It is just part of existing in a real world. What might be an interesting topic for research is how to build a *renewal* process. How do we design for a built-in tropism back towards elegant simplicity[12]?

This is not really network research - it is more about process and production engineering, and probably entails some disciplines we don't even have a good handle on yet (ecology, for example). Reading IEN-1, and looking at Donald Davies seminal work on "The Control of Congestion in Packet-Switching Networks" in 1972, and many other key works that led in to the Internet design, one can have confidence in many of the key features. A radical re-think would almost inevitably lead to many of the same problems after an initial period of even worse chaos. Meanwhile, addressing the current problems through evolutionary (incremental) engineering solutions may well work just fine to stave off collapse.

If long term concerns about the Internet lead to any truly interesting long term research, for me this should be in the area of self-re-enforcing architectures, systems which, like biological entities, have strong tendencies to return to homeostasis. This metaphor should be seen as just that. I am not talking about the operational application of homeostasis literally, as in the use of feedback systems - the Internet already has many of those working quite well nicely, thank you. I am speaking of a meta-design principle - a design pattern for design-resilience. I believe some elements of the Internet architecture already have this property, but it would be good to codify and re-affirm it.

4. CONCLUSIONS

The lessons of the Internet are deep and potentially have wide application. We can apply them in other domains Interesting synergies may emerge.[13]

[11] interface identifiers, part of the end-to-end transport state index, forwarding clues, middle-box state identifiers, access control list keys, etc etc etc

[12] For an elegant elegy on the topic, the IAB plenary talk by Steve Deering at `www.iab.org/documents/docs/hourglass-london-ietf.pdf` is essential reading.

[13] For local UK readers, there are possible paths forward: The future Internet might be a more interesting topic for research if it was to include these parallel studies, and might be a good FISB UK topic. The UK scale is excellent for trying

Such an approach might offer real potential impact for future society, and would really be so much better than the top down command style of programmes being planned and executed across the world today.

One area for foundational work here is indicated, and that is how to deal with multi-dimensional hierarchies. Another one that springs to mind is that of complex systems and emergent properties. I have discussed this elsewhere in the context of the Internet, but it is clear that complex systems are bound up with specifics of a particular domain, and that it is not the least bit clear whether the *stability* or otherwise (for example) of the Internet would translate into a transport or energy system founded on the same principles. There is where I would spend my two cents of tax payer's money on Future Internet.

5. ACKNOWLEDGEMENTS

Thanks are due to G11 and SRG for valuable comments and feedback.

these sorts of ideas. Local advantages might be to achieve the same benefits the UK government has been seeking but without the tendency to centralisation - the database state is unnecessary, and infeasible in any case, and a more nuanced approach might avoid embarrassing errors like NPfIT.