

# Invigorating the Future Internet Debate

Dirk Trossen (Ed.)  
BT Research  
Ipswich, UK  
dirk.trossen@bt.com

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 It has NOT been peer reviewed.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

While many initiatives have been discussing the future of the Internet, the EIFFEL<sup>1</sup> think tank set out to push these discussions forward by assembling a community of researchers debating the potential thrusts towards the Future Internet. This article provides an account of these discussions, being addressed both to the EIFFEL membership and more widely to members of the community interested in medium to long-term developments. We outline immediate problems as well as potentially missed opportunities in the current Internet, while focussing the debate on the need for a Future Internet on the style we conduct research as well as how we design systems at large. Most importantly, we recognize that an inter-disciplinary dialogue is required to formulate actions to be taken and to remove barriers to their realization.

## Categories and Subject Descriptors

C.2.1 [Computer Communication Networks]: Network Architecture and Design

## General Terms

Design, Economics

## Keywords

Future Internet, EIFFEL

## 1. INTRODUCTION

The Internet has become so pervasive and important, that those who do not have ready access to it are deemed to be disadvantaged – on the wrong side of the digital divide. The depth of the embedding of Internet technology within society means that it is now no longer possible to discuss the nature of society without also discussing the role that technology plays in aspects of life that range from cooperative social networking, via competitive business and commerce and information delivery through media organisations, to conflict and conflict resolution - war, politics and human rights. Conversely, whilst it might, even in the relatively recent past, have been possible to discuss the nature of the Internet in abstract technical terms, this is no longer sufficient. Thus, considerations of a Future Internet necessarily encompass the uses and abuses of any proposed technology, and the value judgements that decide which is which.

The technical development of the Internet to date was traditionally achieved through a process of largely free innovation but

more managed community acceptance through the IETF, ISO, ITU-T, etc. Actual deployment has always been partly driven by commercial considerations, but the rate of introduction of technologies that have no extant user base demonstrates that such considerations often extend over at least the medium term. Formally, many of the standardisation bodies are voluntary organisations, with representation drawn from the broad community; the standards do not, therefore, have the force of law (in the main), with the consequence that freedom to innovate is maintained. Conversely, the increasing size of both the established and the likely future deployed base means both that radical innovation and re-architecting is more difficult than it was in the early days of the Internet, and that aggregations around emergent industrial (de facto) standards are becoming increasingly powerful relative to old-style community acceptance mechanisms.

If we learn nothing else from a retrospective view of Internet development, we see that the freedom to innovate has had a profound positive impact on the diversification of the uses of Internet technology, and this, in turn, has meant that it has never been possible to predict the course of such developments. Consequently, whilst it is clear that the present Internet is struggling to cope with current demands, there is a wide diversity of opinions in the large number of 'Future Internet' activities about the likely future and the challenges implicit in reaching it. We do not believe it to be sensible in terms of the lessons learned from the history of the Internet to seek to articulate an ab initio technical design for THE future Internet. Instead, we identify a number of assumptions and questions implicit in the current direction of Internet research, and propose debate at the level of both architecture and the constraints that should be placed on that architecture.

*Future Internet Research is needed.* This seems to be a foregone conclusion, given the numerous initiatives around the globe. But it is unclear whether the term Future Internet research is meaningful in its own right, is a simple statement of the fact that we need to move on from where we are, or is simply a label attached to proposals to secure funding. This article outlines some of the problems of the current Internet that have been recognized by a variety of initiatives. Moreover, it discusses potential missed opportunities, since a myopic focus on problem solving may lead to lost opportunities.

*Future Internet Research is possible.* It is assumed that large-scale Future Internet research programmes are capable of delivering strategic research that addresses concerns beyond those of immediate problem fixing. Such a view relies on the fact that it is possible to identify the most profound research directions, that it is possible to conduct scientifically valid research within them, and that the funding regime encourages the honest exploration of operational boundaries – encourages the possibility of finding that

<sup>1</sup> The EIFFEL Think Tank (<http://www.eiffel-thinktank.eu>) was established in 2008 as a community of Internet researchers and practitioners, debating issues and disagreements surrounding the current and possible future Internet. This article has been produced as a result of the currently ongoing discussions within this think tank.

a proposed approach is unusable. This is difficult, because research funding given in small parcels can be innovative but may suffer from a lack of scale since experiments are expensive to conduct. On the other hand, large parcels of research funding that are capable of conducting large scale experiments (a) can only conduct relatively few such experiments; (b) are too expensive to ‘fail’, so tend to be low-risk and (c) in any case suffer from the fact that the use of the Internet changes once technology is deployed, often in ways that are not predicted by the technology designers. Furthermore, there is growing recognition for the need for theoretical approaches to evaluation [10], such as proofs, theoretical bounds on algorithms or system-level approaches to architecture evaluation that include more than technological aspects, while recognizing that such approaches are likely to remain incomplete with respect to capturing all aspects in a large-scale system.

*How will the Future Internet happen?* Unlike the current Internet, which evolved from a small-scale research network, it is likely that the Future Internet would follow from the large-scale deployed system that permeates many fibres of society and on which people rely. It therefore seems natural to emphasize a focus on the right mechanisms to ensure a continuous and fruitful evolution of this system towards a future one, whilst retaining an ability to sustain and foster innovation. The proper interplay of all interests, i.e., that of researchers, end-users, corporations, community groupings and governments, is crucial for success. This directly affects the style in which we conduct research, posing questions on how to engage with stakeholders and on how to embed processes of market selection and government control into the development of solutions by allowing the expression of concerns in runtime of the system rather than at design time. Equally important, the evaluation of the efficacy or significance of proposed solutions must, in the ultimate, be conducted in the context of the whole stakeholder base with particular attention paid to the potential impact on society.

Future Internet research is more than a funding vehicle. The Internet community has identified immediate problems and challenges that may require to fundamentally re-think the set of mechanisms in use and their architectural parentage. These challenges are rooted in the problems we see appearing in the current Internet. In addition, emerging uses of networks challenge the underlying assumptions inherent in the design of the Internet. Proposed solutions to these challenges can be directly used in the engagement with stakeholders and the establishment of new market opportunities.

Articulating the grand challenges and working towards solutions needs a wider debate as well as concrete work among a growing community of (interdisciplinary) researchers and major stakeholders. The need is clearly understood by all EIFFEL think tank members. Different views exist in respect of what may be missing from the current architecture or why such concepts are missing. This article, among other activities, is not only an attempt to contribute to this debate but also to strengthen activities like EIFFEL that attempt to facilitate such debate and formulate the consent and dissent of such debate for the community to ponder on.

## 2. WHY FUTURE INTERNET RESEARCH?

The Internet has undoubtedly changed the world. It works well—but not perfectly—but well enough most of the time. However, there are recognized problems that cannot be patched within the constraints of the current architecture. Also there are potentially missed opportunities that the Internet architecture is not well

suited to deliver. In brief, the sky is not falling but it is not as high as it could be. It is a pressing question how we could push further in our research. This does not only include the mere need for Future Internet research but also the style of how we perform research.

Whilst we can observe a number of seemingly ad-hoc solutions to the rising number of problems, we believe that the major challenge of the Future Internet is largely architectural and is therefore sufficiently fundamental that it goes beyond a quick succession of patches. We should not imagine we could solve these problems without understanding how a system of such size interacts with the world. The world changes the Internet and the Internet changes the world—society, culture, commerce and technology.

One important observation to make is that, as architects, we should not tip the balance at design time towards the world we want, attempting to bias the Internet towards particular governance and business models. In other words, the architecture must attempt not to prescribe the outcome of tussles in the (future) market place beforehand rather than allow for tussles to commence inside the architecture at runtime. For that, we should provide the minimum substrate that allows the Internet the flexibility to behave in different ways at different times and in different places depending on the outcome of market selection and social regulation [3], whilst retaining levels of performance that render it fit for purpose. Hence we must move from a largely design time to a largely runtime model for resolving potential tussles.

Without the ability to see into the future, there will inevitably be times in which a value judgement with respect to a potential (possibly unidentified) tussle will be determined to be the best approach, avoiding the tussle and choosing a particular path for perhaps technical, moral, ethical, legal, or business reasons. The nature and impact of this choice, however, need to be made explicit as well as understood. Since such choices are inherently constraining, the establishment of an orthodoxy, that results from making constricting decisions, must be balanced by inviting challenge and weighing evidence. For this, it is most important to pay attention to the evolutionary mechanisms of the Internet—the aspects that determine how evolution progresses and, indeed, whether evolution progresses at all. Decisions made at this point must remain relevant and fresh for at least as long as the current Internet has proved valuable, in a world in which Moore’s law continues to apply. Investment of time, effort and hard cash in widespread changes to the whole system will not occur unless such changes both deliver in the timescales needed for cost recovery and continue to give returns over many decades in an constantly evolving technological, economic and societal environment. Over such timescales, we cannot and should not predict what the world will be like. Paying careful attention to catering for evolvability is therefore the only reasonable approach. However, the lack of predictability might also question the suitability of the platform itself to evolve at required speeds. Hence, more radical changes, or larger evolution step size, might need to be necessary to ensure overall platform evolution.

Each of these aspects is discussed below in a little more detail: immediate problems, potential missed opportunities and evolutionary mechanisms. We then argue that a different research culture is needed to address these problems. For this discussion, we define the scope of Future Internet research as research on all the generic aspects of the Internet. This is infrastructure but not just ‘provided’ infrastructure: it includes the protocol handlers and interfaces on end systems. We also include infrastructure that may

not strictly be part of the Internet, but provides services that are often needed to use the Internet, such as naming and discovery services.

## 2.1 Immediate Problems

Over the years, some major problems have proved amenable to fixes within the architecture, e.g., inter-domain differentiated services or classless inter-domain routing. However, other problems—mostly to do with control and management—have proved resistant to quick fixes. Some remain as major outstanding problems. Others have been fixed in ad hoc ways by violating the architecture, which has resulted in a mess that limits future evolution. Some examples are:

*Resilience, failure tracking & management:* The Internet design is renowned for its robustness to failure. Indeed failures often do heal automatically, but not quickly. Also many failures are not amenable to automatic solution, e.g., due to human errors in configuration. It is generally believed that the Internet of today does not have effective solutions to these problems, although there is no consensus on what needs to be changed or whether such solutions would be necessary on the architectural level, in particular when it comes to the necessity to evolve from current approaches.

*Availability & robustness to attack:* The Internet is continually being used as the means for malware to attack both services and the Internet infrastructure itself. Solutions to these problems often block innovative legitimate uses of the Internet as well as illegitimate ones, effectively slowing down the Internet's evolvability. Proper architectural support to address the root means of these attacks is needed, but there is no consensus between the contending partial solutions. Therefore, as well as fixing known inherent problems, any solution will also need to evolve to address yet unknown forms of attack—but still without limiting evolvability. For this, perhaps, it is desired to construct an incentive structure that would encourage legal (and “good” by some metric) behaviour in using the facilities available. Furthermore, a potentially radical re-thinking of sender-oriented paradigms like today's IP is required, shifting the balance of power towards the receiver of data and information.

*Information security scalability:* The state of the art in information security techniques is sufficiently robust to assure any form of security, except that the techniques do not scale to global proportions in non-hierarchical groups. The Internet is designed so that information security can be built over it, end-to-end. So this is seemingly not an Internet architecture problem although a holistic view on information security would demand architectural solutions since the lack of globally scalable information security techniques does hold back evolution of new developments over the Internet. For instance, those in which every client's identity, not just service identities, needs to be represented in the system often require proven identities over an interlinked social network rather than a hierarchical organisational one (c.f. the evolution of mobile IPv6). Another aspect of information security is that of information accountability. While the Internet can cause information to be shared or not, once it has been shared at all, any control is essentially lost of any further sharing and exposure and are dependent on some vague sense of trust in those with whom we have shared.

*Scaling for more extreme dynamics:* The dynamic range of the Internet architecture is hitting its limits. For instance, increasingly the inter-domain routing system cannot converge quickly enough following a change, leaving longer periods of disconnection. More sites are connecting to the Internet through multiple links to

improve resilience, but the inter-domain routing system is designed so it then has to treat these sites as distinct networks rather than as stubs off a single-provider network. This makes the routing system appear much larger without the Internet growing at all. Also the Internet's congestion control mechanisms have hit the end of their dynamic range due to higher bit-rates.

*Resource accountability:* The Internet architecture allows everyone to use any resource anywhere on the Internet to the extent that they want. However, at present, network operators are deploying boxes to limit or block communication with certain users or by certain applications, in particular those posing an economic threat to the provider. The response from application developers has evolved into an ‘arms race’ with more and more checks and blocks being placed in the way of truly innovative applications. It is now recognised that end-to-end transport protocols cannot and should not be expected to share out capacity. But the Internet architecture still contains the inherent assumption that they do. Even if networks were trying to share out capacity without making judgements about content, the architecture does not reveal the information they need to make other networks and their users accountable when they are over-using stressed resources. The consequent inability to properly limit free-riding (or to deliberately allow it) leads to uncertainty over whether capacity investments can be recouped, which in turn negatively affects the whole value chain of the Internet.

*Network-application coordination:* Over the years, the application programming interfaces (APIs) at the top of the TCP/IP protocol suite have become ossified and stale, but more importantly they have become almost impenetrable. In the downward direction, middleboxes (e.g., firewalls and network address translators) largely recognise only those protocols that existed when they were deployed. So they block out all attempts by applications to use new APIs to new lower layer protocols and services. In the upward direction, applications cannot find out about the network or their paths through the network in order to create richer services themselves—services that could exploit knowledge of network topology, network failures or traffic characteristics. Evolution of capabilities is a real issue here. However, such an evolution is likely to be only accretion. Eliminating capabilities that were once there often has the negative side effect of abandoning the users of those capabilities. So, the pressure is always only to add, not change. This, in turn argues to some extent that evolution in the core may not be such a great idea, but rather getting the core “right enough” and “flexible enough” that it has a very long lifetime.

## 2.2 Potentially Missed Opportunities

The following list is not meant to be representative or exhaustive. It merely gives known examples of big new opportunities that the Internet architecture might stumble over. The intent is to show that a more generic (or less monolithic) architecture would be useful:

*Interconnecting the information & physical worlds:* The vision of ubiquitous computing [Weiser91], where devices merge into the background fabric of life is only partially realised at present. Whilst there are vast numbers of embedded computing devices involved in many of the aspects of our lives, most of these are not currently networked, though increasing integration shows that this is practicable. The Internet architecture has been successful for non-embedded devices, but there are clearly problems that arise in having an architecture based on end-to-end principles working with endpoints that are resource constrained to the extent that they

cannot protect themselves, nor rely on boundary protection, nor even necessarily support traditional IP protocols, either for reasons of constraint in the low bandwidth connectivity they may have or because the IP addressing modes are inadequate to capture the more data-oriented, partially connected nature of the system. It remains unclear whether there is a sound business case behind the broad vision of pervasive computing or whether architectural conflicts are preventing or slowing the take-up of such technologies. Whatever the reason might be, further research is required to understand the cause and the potentially missed opportunities.

*Natural social interaction:* Over the last two decades human communities have exhibited a strong trend towards more continuous remote as well as local contact with each other. But despite young people staying in regular touch and online meetings having become a fact of work life, the technologies used are still far from natural, relaxed social interaction. Again, these problems may not be rooted in anything architectural, but one does wonder why such obviously valuable capabilities have not emerged.

*Governance models:* The current debate over the future Internet governance in the area of name spaces and address allocation has been characterized as a choice between unilateral and multilateral government control, which is a false dichotomy. There is a third choice – no centralized management – that allows maximum participation, nationally and regionally. It should not be characterized as preventing any entity from participating in the governance of Internet.

*Cater to new communication paradigms:* The internetworking solution of the current Internet is based on a simple packet-based forwarding paradigm between endpoints known to each other. This paradigm directly stemmed from the predominant telephony model at the time of creating the Internet, transferred onto a world where devices and machines interact. But new paradigms of communication are appearing, e.g., in the sensor world (diffusion models) or in the content world (P2P, publish-subscribe instead of send-receive, information labels instead of endpoint addresses, IP multicast based on delivery groups instead of endpoint addresses). Overlaying on top of IP allows for integrating some of these paradigms, often leading to inefficient usage of resources on the IP layer and below. Support for a more polymorphic nature of networking, natively supporting a variety of paradigms, holds promises for the rise of new services and applications as well as more efficient use of resources.

The Internet architecture is often likened to an hourglass, with the internetwork layer at the neck enabling a wide range of applications above the neck to use a wide range of connectivity technologies below the neck. In addition to missing opportunities for new uses of the Internet—above the neck of the hourglass—the Internet architecture today only exposes the lowest common denominator of the capabilities below the neck. Therefore, as technology advances, the architecture is possibly not best placed to take advantage of these developments. For instance, memory and storage are now considerably cheaper in relative terms than in the Internet's formative years. Does the current architecture allow us to exploit such shifts in the original assumptions to their fullest extent, or does it hold us back?

### 2.3 Evolutionary Mechanisms

When considering a large-scale system such as the Internet, the following observations can be made with respect to evolvability. First, *there is a need for evolution* as a gradually developing process, like for any large-scale system. This system evolution is par-

ticularly important considering the evolution of society due to the impact of the very system itself. In order to understand the suitability of the system to evolve, we need to understand the dynamics forcing the changes and devise an architecture that is suited for these dynamics to commence in runtime. Second, the *scope of the dynamics* effecting change of the Internet is widening. The Internet has become more than a technical artefact – it has transformed from a network for geeks to a crucial infrastructure for society and business. Its impact on these areas is obvious, from e-commerce to e-government, the change in the perception of privacy to many other societal changes since its introduction. The virtual and the real world abide by similar rules regarding human rights and respect for personal space as guiding principles. Hence, the question of evolving the Internet is not a mere technical one anymore. Third, *evolution speed is increasing* with the advances of technology. For instance, memory is becoming so cheap that solutions for caching vast amounts of content locally are likely to transform the way we deal with content. While the dynamics of an industry in which functions (and control) can be shifted in real-time still need to be understood, such increasing speed of dynamics is well observed, e.g., in [10].

The Internet's ability to evolve becomes compromised when the architecture does not allow legitimate concerns to be expressed after its original design. As a result, people solve their problems in ad hoc ways, adding carbuncles in violation of the original architecture. Then subsequent requirements are even more difficult to satisfy, because of all the feature interactions with the exceptions to the original architecture. The root of this problem lies deep in the processes that we use to design architectures. Much emphasis is placed on the *design phase* of the architecture, with requirements phases and use case definitions, accompanied by processes of standardization. This inevitably leads to an emphasis of the concerns that are important to the players who are deeply involved in this phase while often neglecting the concerns of the actors entering the scene after the solution has been fixed. Hence, the outcome of the design is often seen as a fixed architecture, as opposed to explicitly aiming at designing an architecture that can fulfil evolving requirements.

This Newtonian-Cartesian approach towards system design assumes the ability to capture all relevant concerns and therefore resolve the most probable run-time tussles at design time. The widening scope of the Internet beyond mere technology and the observed increase of ad-hoc solutions to concerns of actors after the design of the original architecture bring this design process into question. The authors of [9] describe a shift from these reductionist Newtonian-Cartesian towards Darwinian approaches, where the *evolutionary kernel* (i.e., a component that has proven successful for multiple uses so it will act as a platform for evolution around it, see [5],) becomes the design process itself, i.e., a process in which concerns of actors are incorporated into the system at runtime, e.g., through mediating or isolating conflicting concerns, recognizing the inability to cater to all possible requirements during design time.

Such a shift in processes requires an understanding of what had been good and what needs to change to cater to the future. For instance, the Internet's hourglass model has a relatively successful track record of innovations above and below it. But the neck of the hourglass itself has proved highly resistant to change, potentially constituting an evolutionary kernel. It is important to understand that there is potentially not just one such kernel. There also needs to be innovation in the surroundings of Internet service provision. For instance, the contractual interfaces between cus-

tomers and providers are hard to change (access control, authentication, etc), as are the contractual interfaces between networks (border gateway protocol, key exchange, etc). Hence, programmability on all architecture levels becomes key.

In our rush for change however, we must not lose what was good about the previous architecture—we must be able to recognise a baby in the bath water. No good reason has been articulated for why all the original design principles [2] should not still stand. However, it is recognised that the end-to-end design principle is hedged around with stronger caveats than before. There have been some attempts at defining additional principles, and better articulating the concerns around the end-to-end principle (see for instance [1, 7, 4]). But we must be open to understanding how we can establish design processes that allow evolution towards the future requirements without delaying the progress of that evolution unnecessarily. Hence, we must strive for a culture of design that results in designs that are founded on the (empirically successful) past but that nevertheless allows for a future that is characterised by change.

### 3. RESEARCH STYLE

The basis of scientific research lies in the ability to formulate and test falsifiable hypotheses. The role of engineering is to create, evolve and maintain operational systems according to a particular design brief. The Internet provides an environment that is rich in possibilities for research that is experimental and analytical, which, at the same time, must be set in the context of engineering; likewise, Internet engineering must respect the need to use the engineered system as an experimental platform and as a platform for innovation, both of which might cause the underlying design brief to change.

The Future Internet is, consequently, more about process than product. Although it is likely that a Future Internet will result from agreement by committees representing industry players and governments, it is crucial for individuals (including researchers) to understand how to influence the key decision makers to eventually adopt the ‘right’ solutions. The present Internet overtook the plans of committees, which have historically focussed more on engineering than innovative vision and would probably have adopted Broadband ISDN, for example. Some believe the future Internet will come about through the same institutions that fostered the current Internet; the networking research community and the Internet Engineering Task Force (IETF), but this is unlikely to be realistic, simply because the importance of the Internet has changed with time and the list of stakeholders with an interest in design outcomes has grown. As things stand, it is undoubtedly the case that many proposals will be standardised in a variety of committees. Those proposals that are most worthy and manage to attract support of key stakeholders will be deployed, and those that survive the rigours of the marketplace will become the Future Internet.

As a result, we believe it is likely that the Future Internet will result from a process of community innovation with market selection, albeit with a broad community who have an increasingly heterogeneous, contradictory and sometimes irreconcilable set of perspectives and success criteria. Nevertheless, we believe it to be important to the success of such endeavours to recognise the importance of the duality of *free innovation* and *solid engineering* and hence to regard support for both of these as a critical input to the decision making process that controls Internet evolution. Thus we need input from *those with a reductionist viewpoint*. The research community that brought the Internet through its childhood

has been characterised by innovative engineering based on an appeal to reason and deduction, with just sufficient gloss of scientific rigor and projected utility to secure venture funding for the idea of the moment. This is what is needed to develop new paradigms that address ‘Missed Opportunities’, but it is much less appropriate for solving the ‘Immediate Problems’ or for working on the ‘Evolutionary Mechanisms’.

But we will also need input from *those with an empiricist viewpoint*, i.e., the engineering community whose primary focus lies in analysing and fixing a pre-existing living, evolving, global-scale system, which requires somewhat different skills and organisation from the process of inventing something new and relatively independent. Analysing the root cause of known Internet problems, and predicting the effect of changes to the underlying working system is a ‘Big Science’ (but also a policy) challenge. It requires rigorous experimental method, sharing of data, duplication of results, integration of observational (e.g., in social science) as well as explorative (e.g., in engineering research) styles of research, and so forth. It also requires long-lived, often multi-disciplinary, research teams, and long-lived institutions to foster consensus and provide continuity. And last but not least, it not only requires existing stakeholders to be consulted, those with commercial interests and those with operational experience, but also to enable potential future stakeholders having a voice in the process. Of course, it also requires engineering innovation, but that will be a small part of a much larger research programme to assess new ideas and fit them in to the wider picture.

Finally, any who believe that defining the Future Internet is a largely technical problem are also unrealistic. Research teams will need to be multi-disciplinary, with experts in different cultures, in national and international law, in economics, in human rights, etc., not only experts in inter-domain policy-routing or network control theory. If evolutionary mechanisms are to shift to centre stage, engineers have to understand which aspects they should not pre-judge by hard-coding them into the architecture. Then the Internet can adapt to the world at run-time, rather than the world having to adapt to what was embedded at design time.

### 4. TECHNOLOGY & SOCIO-ECONOMIC CHALLENGES

A focus of our discussions was placed on the areas of technology, economics and society. The following section gives a brief insight into these discussions.

#### 4.1 Technology: Looking Backward & Forward

How do we find out the ‘best’ network design? We can look back to the early days of networking for clues. Around the late 60’s and early 70’s, the network and computing research community as a whole designed, implemented, and operated ARPAnet first hand. A number of basic design principles were derived from the process. Lessons were learned from that design and operation experience; the end-to-end principle became fundamental in the Internet architecture. The hourglass model emerged as a result from meeting the need to support diverse higher layer protocols and applications, while bridging different underlying communication technologies and enduring their changes. The Internet is an artefact that can and should be studied just as physicists study the world. The fact that we designed a network does not mean that we can fully understand how it works. We build networks by specifying individual components and we specify a protocol in a static way, which does not lead to a description of the dynamic, collective

behaviour of multiple components when they interact through protocols. Furthermore, multiple (in fact, a large number of) protocols exist in the system and we do not know how to derive the combined behaviour out of the interactions among various pieces.

Thus after (or even in parallel to) the engineering phase, we as network researchers have to take on a scientist role to understand the world we just built. We need to rigorously follow a set of established scientific methodologies to guide our process and procedures in exploring and understanding the results of our own design. This includes examining (and hopefully understanding) the emerging behaviours that we probably could have not understood before the design. It also includes the ability to properly observe the phenomena, which occur in daily operation. For that, an inherent ability is required to measure data in the system. Judging the relevance of data to be measured, however, is often a challenge, given the (often) unclear nature of (future) observations. However, mastering this challenge might hold a valuable price that sets us apart from natural science, namely because we created artefacts in the first place, the understanding we gain from its operation and experience can be fed back into designs of new artefacts or tunings of existing ones. Hence, we first design systems as engineers. We then study the artefacts as scientists. Our scientific findings are then feed back into the design or revision phase. This cycle continues and never ends.

#### 4.1.1 A Few Lessons Learned

We need a full understanding of the driving forces behind the Internet's success. The Internet would not have succeeded so greatly without Moore's Law. Computing technologies are moving forward with accelerated speed. The Internet architecture facilitated the technology advances. The rapidly advancing technologies in turn drive new application developments and user population growth on the Internet.

Technology advances and Internet growth also created new demands on the architecture. The need for security, manageability, and scalability showed up over time. Today they are more pressing than ever, as they were not promptly identified and fixed ahead of the crisis. We need to continuously identify and address these new demands. As best stated by Feynman "what is the source of knowledge? Where do the laws that are to be tested come from? Experiment, itself, helps to produce these laws, in the sense that it gives us hints." That is, one must stay at the frontier of the practice to gather the hints for new problems and new demands.

An unfortunate fact is that there has been a big gap between reality and the research community's understanding of it. Since the Internet commercialization in mid 90's, the networking research community gradually lost touch with the frontier of the Internet, hence lost the opportunity to observe real problems first hand. The community by and large retreated back to work on isolated or point problems, and used simulations or small, isolated test beds at best for design evaluations. The research community's lack of attention does not mean real problems do not occur, but only that the problems are solved by others.

There is also a lack of understanding that successful network architectures should change over time. All new systems start small. Once successful, they grow larger. The growth will bring the system to a new environment that the original designers may not have envisioned, together with new requirements that must be met. For example the security threats facing the Internet in recent years should not be blamed upon the inadequate design of the original architecture. Rather, it is due to poor understanding of its

limitations. Continued success requires continued scientific research on networking practice, to identify new problems and evolve the architecture to meet the new demands. And we need to strive for clear evidence on such successful networking practices.

Last, but not the least, we need to develop a good understanding of networking architecture advances. It is often stated that the Internet architecture has not been changed over the last 30 years (although depending how one defines changes, this statement may or may not be entirely true in reality [8]). But how should changes be made? Would it arrive as a revolution, or should the adjustments being made continuously as in an evolutionary process? In any case, the Internet architecture should be prepared to evolve either based on smooth evolution or based on rapid introduction of fundamental paradigms.

## 4.2 Economics: Design for Sustainability or Sustaining our Current Design?

The Internet was designed as a packet routing fabric that places the ends in the centre with an assumption of a dumb network. This has led to a cycle of innovation that is largely driven by this edge. But this focus on edge innovation has left many problems unsolved, such as QoS, multicast and others - problems that require coordination among stakeholders. Considerations for deployability and economic sustainability are crucial. This requires a dialogue between communities that have mostly been isolated in their work, namely technologists and economists. As Dave Clark puts it: "We do not know how to route money". It expresses a sentiment in the technology community to be either ignorant or weak when it comes to economic questions. But there is a similar ignorance in the economics community towards questions of technology design. Hence, little work has been sprouting from the intersection of technology and economics, although this is changing slowly. This goes beyond the problem of merely identifying the right people to involve. It also needs an approach to overcome the existing isolation between the 'camps' that is manifested in language, approaches and techniques. We believe that only a clear identification of (some of) the challenges that will require the joint dialogue will motivate such common debate. A few of these grand challenges are outlined in the following.

*Joint System Design:* Designing a technical system creates an economic one, while the latter is enabled by a variety of technical systems. In reality however, the process of (technical) system design is mostly disjoint from the process of designing business models and strategies for sustaining them over a period of time. Combining these two processes is difficult, largely because of the communities that are required to interact. But it is hard to see how challenges on, e.g., sustainability of systems, can be solved without such joint design process. A solution to this problem will not only have an impact on the design of systems but also, for instance, on the way we educate talent in their understanding of these fields, as has been recognized in, e.g., [6]. For this to happen, however, we need to accommodate the differences of research styles that exist between research fields, like economics, engineering and social science.

*Sustainable Value Chains:* It is understood that value chains change with innovation through the entry and exit of market players. But in order to stimulate any innovation, a certain degree of sustainability is required to make investment happen and maintain a healthy value chain throughout the expected timeframe for the return of this investment. In order to sustain such healthy value chains, it is crucial to account for the usage of resources, in particular for the side effects of that usage, such as congestion. This

requires mechanisms and overall architectural approaches for a resource accountability framework - a problem directly relying on some solution to the joint system design challenge.

*Validity of Current Business Structures:* The particular (technical) approach to the Internet has created business structures that evolved around it, such as expressed in transit and peering relationships of autonomous systems. Any evolution of technologies, in particular any fundamental change to the current Internet, will undoubtedly have an impact on these existing structures. Too radical a change will cause problems in adopting the change – and the mere lack of understanding the proper impact can be as much of a showstopper. Hence, technical and economic migration strategies from are crucial for a wide adoption of proposed changes.

### 4.3 Society: The Dual Roles of the Future Internet and Social Organisation

As mentioned earlier, the Internet is both a reflection of parts of our social structures and has come to have a strong influence on our social structures. From early days it was a vehicle for both enabling email based communication and for simply improving the information flow between parties which would have otherwise exchanged the information, but more slowly or in lesser quantities. In addition, along with this beneficial relationship with social structures, hand-in-hand came anti-social opportunities. Birth/death records, medical records, banking records and so forth were kept long before there was an Internet, but the Internet not only made them aggregatable, but also made it simpler for malefactors to get at them. It was only as the infrastructure became increasingly integrated into, and critical to, our society that attacking it also became increasingly worthwhile.

It is important to understand that as a reflection of society, the Internet is always a partial reflection. As such, if we consider that it will evolve to provide increasing aspects of social infrastructure requirements, we are unlikely to be accurate in predicting where the next step will be. In fact, some of the innovation comes from other quarters. Who would have thought that carrying around small wireless cell phones with tiny keyboards would turn into instant messaging and from there make the leap to the Internet and soon into all the different modes of social networking? Hence, innovation will always have an element of surprise for some stakeholders.

The demands of improving social communication and reflecting social structure, while at the same time increasingly addressing issues of privacy and safety in a connected world, could follow many possible alternative paths. As we consider a future architecture for the Internet, we must recognize that the revolutions that have been reflected in it were mostly not of the making of the researchers, but of people solving particular problems for other people in social structures. The vast majority of those efforts did not cause significant disruption or revolution. It is most important to note that we are notoriously bad at predicting which ones will remain innocuous or die and which will be transformative.

One of the interesting social challenges we are faced with is to deal increasingly with overload. There is too much information. There are too many services that want to claim our trust. There are too many options and too many individuals who want our attention. A challenge will be to evolve approaches that reflect the human and social approaches to dealing with overload. We are already doing that in what are probably simple ways in social networking contexts. We group our friends; we create channels for topics, and so forth. We are beginning to cluster the world around us, but we are only at early stages. Newspapers were a mechanism for filtering,

organizing and limiting information that otherwise would overwhelm us. With the demise of newspapers, what elements of the almost infinite flow of bits will bring order that is reflective of the human mind and human social structure? In the longer run, will that also allow each of us to retain a somewhat personal view in large social structures? How will our individuality and privacy be retained?

An interesting societal question arises around is the issue of whether governance has an impact on the Internet or the Internet on governance. One can also turn these questions around. It is clear that the low-cost and pervasive availability of a uniform communications substrate has had an immeasurable impact on our global society. Historically explorers circled the world and laid claim on behalf of their home countries to other lands, thus beginning the political and economic connectedness around the globe. The presence of the Internet has qualitatively changed the nature and degree of that connectedness. In the current economic and political situation, no country can make decisions with a mere local effect. There is no more isolation. Given that, one must consider the relationship between the Internet and governance. And perhaps even more importantly, the Internet may change forever governance of, by or for a people. Blogging and cell phone cameras are having profound effects on the capability of individuals to influence their governments at times when the governments may not want that. Even more, there is an interesting set of questions around governance and perhaps the meaning of democracy when it can possibly be fully participatory, rather than only representative. This is likely to have an impact on, e.g., regulation when considering a growing role of end users in the participation of the Internet, i.e., end users potentially grow into an essential part of the Future Internet, moving away from their current pre-dominant role of a mere consumer. How this will affect ways to regulate certain parts of the Internet will be important to understand.

As we examine or enable evolution in the Internet, there may be duelling forces, one to enable new capabilities as demanded to provide increasingly functional infrastructure to our societies and the other to understand the social ramifications of changing the infrastructure.

## 5. CONTINUING THE DISCUSSION

The discussions as being summarized in this article have proven to be useful to a community of individuals, a community that meets regularly for deepening our understanding of these issues. But it has been felt that the debate needs a wider involvement of the community and a better stage than physical meetings. To begin with, it was concluded that three things could benefit from a community-wide effort. First, a *repository* of relevant work in the Internet area, enabling proper referencing but also discussion on its relevance. Second, an *agreement on disagreements*, including the precise nature of the disagreement, is crucial for driving our debate forward. Such disagreements, e.g., around the relevance of the end-to-end principle or around the importance of areas like virtualization, are driven by formulated viewpoints but also discussions on these viewpoints. This, hopefully, will lead to a taxonomy of areas that need to be addressed in the wider community in order to push the Future Internet significantly forward. Last but not least, a *forum for architectural debate* has been identified as being necessary. Such forum should allow for debating deep architectural issues that are otherwise hard to publish in common conference venues. As the platform for these three aspects, an online one has been created, decoupling the debate from the direct need for physical meetings (although the notion of the think tank as a driving contributor to this debate will continue).

For this purpose, the FIPedia site has been created at <http://fipedia.org/fipedia>. This article is, among other things, an invitation to the community to join our debate on this site.

## ACKNOWLEDGEMENTS

The ideas in this article reflect numerous discussions among participants in the EIFFEL think tank. Opinions expressed are drawn from consensus views among the individuals participating in this initiative but cannot be represented as the exact opinion of any individual. There are undoubtedly more people that have contributed ideas verbally or electronically and we acknowledge all these contributions.

## REFERENCES

- [1] B. Ahlgren, M. Brunner, L. Eggert, R. Hancock, S. Schmid, “Invariants: A New Design Methodology for Network Architectures”, Proceedings of SIGCOMM 2004 Workshop on Future Directions in Network Architecture (FDNA 04), 30 Aug 2004, Portland, Oregon, USA.
- [2] D. Clark, “The design philosophy of the DARPA internet protocols”, ACM SIGCOMM’88, Computer Communication Review 18 (4) pp. 106–114 (August, 1988).
- [3] D. Clark, K. Sollins, J. Wroclawski, R. Braden, “Tussle in Cyberspace: Defining Tomorrow’s Internet”, IEEE/ACM Transactions on Networking 13 (3) pp. 462–475, June 2005
- [4] D. Clark, M. Blumenthal, “The End-to-End Argument and Application Design: The Role of Trust”, Proceedings of Conference on Communication, Information, and Internet Policy (TPRC), 2007
- [5] C. Dovrolis, “What would Darwin think about clean-slate architectures?”, ACM SIGCOMM Computer Communication Review 38 (1) pp. 29–34, 2008
- [6] P. Mahonen (ed), D. Trossen, D. Papadimitrou, G. Polyzos, D. Kennedy, “Future Networked Society”, EIFFEL whitepaper, 2006
- [7] A. Ford, P. Eardley, B. van Schewick, “New Design Principles for the Internet,” Proceedings of Int’l Workshop on the Network of the Future (Future-Net’09), June 2009
- [8] M. Handley, “Why the Internet Only Just Works,” BT Technology Journal Vol 24, No 3, July 2006
- [9] R. Hollingsworth, K. Müller, “Transforming socio-economics with a new epistemology”, Socio-Economic Review, vol. 6, pp. 395–426, 2008
- [10] Network Science and Engineering Council, “Network Science and Engineering (NetSE) Research Agenda”, at <http://www.cra.org/ccc/docs/NetSE-Research-Agenda.pdf>, 2009
- [11] D. Trossen (ed), “The Core-Edge Story”, whitepaper of the Core-Edge Dynamics working group at the Communications Futures Program, MIT, 2005.