Future Internet

EIFFEL Report:

Starting the Discussion

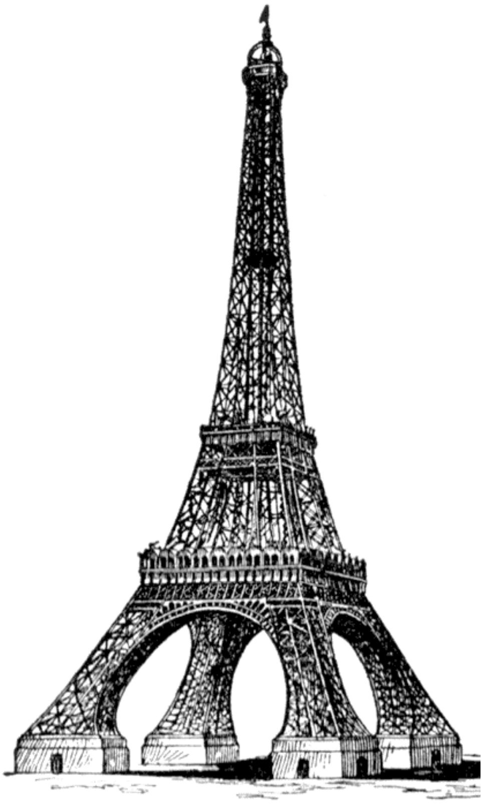

Picture Source: FCIT

FINAL VERSION. 13TH OF JULY 2009

The purpose of this document is to provide a first account of the discussions within the EIFFEL think tank, set within the context of broad commercial, regulatory and academic concerns about the nature of and process of arriving at the Future Internet. It is, therefore, divided into three parts: technology, business and society.

The report is addressed both to the EIFFEL membership and more widely to members of the community interested in networking developments in the medium to long term. Whilst the report is not intended as a verbatim account of any specific Think Tank meeting, it is founded on the discussions that were held within the first two meetings.

The overall aim of the report is to stimulate focussed debate on desirable research and regulatory goals, actions to be taken in reaching them, and potential barriers to their realisation.

Authors: Dirk Trossen (ed), Bob Briscoe, Petri Mahonen, Karen Sollins, Lixia Zhang, Paulo Mendes, Stephen Hailes, Borka Jerman-Blaciz, Dimitri Papadimitrou

EXECUTIVE SUMMARY

"THE EIFFEL Report"
Starting the Discussions

The Internet has become so pervasive and important, certainly in Western society and increasingly elsewhere, 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 accurately to predict the course of such developments. Consequently, whilst it is clear that, even in purely technical terms, the present Internet is struggling to cope with the current demands placed upon it, there is a wide diversity of opinions (of which none may be correct) in the large number of 'Future Internet' activities around the world about the likely future and the challenges implicit in reaching it.

So far as the EIFFEL project is concerned, 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.

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

Assumption: *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 (likelihood) of finding that a proposed approach is unusable in practice. This is difficult, because research funding given in small parcels can be innovative but may suffer from a lack of results in scale since such experiments are expensive to conduct, and 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.

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 already 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 existing 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, corporations, community groupings and governments, is crucial for success. This directly affects the style we as researchers use to 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 such concerns in the runtime of the system rather than at design time. Equally importantly, 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 impact it may have on society.

The grand challenges that need to be addressed. Research for the future Internet is more than a vehicle for obtaining funding for the next five years. The Internet community has identified both immediate problems and challenges that may require a fundamental rethinking of the set of mechanisms in use and their architectural parentage. These challenges are rooted in the problems we see appearing in the current Internet, such as the vulnerability to attacks, scaling for more extreme dynamics or the lack of resource accountability. However, in addition, emerging uses of networks challenge the underlying assumptions inherent in the design of mechanisms for the Internet - for example, the growing integration of physical and digital world or the rise of new communication paradigms. Proposed solutions to these strategic challenges can be directly used in the engagement with major stakeholders and the establishment of new market opportunities.

Socio-economic understanding is key. Given the scope of the Future Internet, we need to address the socio-economic challenges that the Internet and its evolution (will) pose. This includes the place of the Internet: in an individual society under government – including issues of control, monitoring, and censorship; in a global collection of societies – including issues of ownership, regulation and trade; and in the value chains of current and future industry. We also need to identify new trends that will have impact on established structures, e.g., in regulation. Such understanding is particularly important when designing architectures and systems, which necessitate a growing strand of research on the Future Internet in the light of an understanding of how the technologies we create will shape future society. This will require increased dialogue between communities that have had little exchange in the past as well as a mutual perception of the value of different perspectives.

There is need for a wider debate. 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 report, 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.

WHY FUTURE INTERNET RESEARCH?

The Internet is an amazing artifact that has changed the world and is continuing to surprise us. It generally works well—not perfectly—but probably well enough most of the time. However, there are some recognized **immediate problems** with the Internet 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. This section intends to shed light on both issues.

Whilst we can observe a rising number of seemingly ad-hoc solutions to the rising number of problems, we strongly 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, however, imagine we can solve these structural and architectural problems without understanding how a system with the size of the Internet 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 one particular model of governance and business model. In other words, the architecture must attempt not to prescribe the outcome of particular **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 mechanisms [Clark05], 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 a constricting decision 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 and we seek to satisfy the demands of an industry with the clock-speed of a fruit fly. 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 a 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 problems of this magnitude. In this report, 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 liberally include infrastructure that may not strictly be part of the Internet, but provides services that are often needed to use the Internet, such as key management infrastructure, naming and discovery services.

Immediate Problems

Over the years, some major problems have proved amenable to fixes within the architecture, e.g., inter-domain differentiated services (Diffserv) or classless inter-domain routing (CIDR). However, other problems—mostly to do with control and management—have proved resistant to any quick fix. 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 of these are:

- **Resilience, failure tracking & management**: The Internet's distributed design is popularly renowned for its robustness to failure. Indeed failures often do heal automatically, but not quickly. The result is an increasingly unreliable service. Also many failures are not amenable to automatic solution, being due to human errors in configuration and so forth. It is generally believed that the Internet as of today does not have effective solutions these problems, although there is no consensus on what needs to be changed or whether such solutions would even be necessary on 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 a 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 for 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.

- **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. Without architectural support, their attempts to identify free-riding are largely arbitrary. The response from application developers has developed 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) only recognise 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 life-time.

- **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 disconnectivity. 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 since higher bit-rates require higher accelerations to reach them.

Potentially missed opportunities

The following list is not meant to be representative or exhaustive. It merely gives known examples of big new ideas 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 through applications such as IP and mobile telephony, mobile data or sophisticated media players 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, other than for niche applications, 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. Furthermore, we need to understand the societal impact of embedded technology, potentially affecting the lives of every human being (or selected groups).

- **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 regular online meetings having become a fact of work life, the technologies used are still primitive and limiting—still far from natural, relaxed social interaction. Text interfaces on mobile phones, tiny jerky video images or crackly audio conferences with background interference and clunky or non-existent session initiation and floor control. Again, these problems may not be rooted in anything architectural, but one does wonder why such obviously valuable capabilities haven't 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 of having no centralized management that allows for maximum participation, nationally and regionally. It should not be characterized as preventing any entity from participating in the governance of Internet. Many world-wide models of international organizations have shown successful coordination in billion-dollar industries such as SWIFT, FIFA, IATA etc.

- **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. A 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 hour glass, with the internetwork layer protocols 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 like those mentioned above for new uses of the Internet—above the neck of the hour-glass—the Internet architecture today only exposes the lowest common denominator of the capabilities below the neck. Therefore, as technology advances, the architecture is not necessarily 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 exploiting such shifts in the original assumptions to their fullest extent, or does it hold us back?

- **Broadcast esp. radio link technologies**: Although multipoint and broadcast links have been part of the Internet landscape from the start, their useful characteristics tend to be suppressed and their pathological characteristics (interference, fading etc.) are often brushed to one side. Essentially, early on, the Internet was optimised for point-to-point wireline links.

Evolutionary Mechanisms

When considering a large-scale system as the Internet, the following observations can be made with respect to evolvability:

- There is a need for **evolution as a gradually developing process**, like for any large-scale system. This evolution of the system is particularly 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. These dynamics will need to define the required stepsize in evolution that is being necessary and therefore the changes in the underlying architecture that are being required.

- The **scope of the dynamics** affecting change of the Internet is widening. The Internet has become more than a technical artifact – it has transformed from a network for geeks to a crucial infrastructure used in 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 to 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.

- Evolution **speed** is increasing with the advances of technology. For instance, memory is becoming so cheap, in particular compared to the formative years of the Internet, that solutions for caching vast amounts of content locally is 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., [Tro05]).

We have already seen that the Internet's evolvability 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 plethora of exceptions to the original architecture.

The root of this problem lies deep in the **processes** that we use to design architectures and solutions. 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.* This Newtonian-Descartian concept of system design, relying on such requirements and use case definition phases, 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. [Hol08] describes a shift from these reductionist Newtonian-Descartian towards Darwinian approaches [Dovrolis08], 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 [Dovrolis08],) becomes the design process itself, i.e., a process in which concerns of actors are incorporated into the system at runtime, recognizing the inability to cater to all possible requirements during design time.

Such a shift in processes towards a Darwinian approach requires an understanding of what had been good in the current architecture and what needs to change to cater to the future. For instance, the Internet's hour glass model has a relatively successful track record of innovations above and below it. But the neck of the hour glass 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, even the contractual interfaces between customers 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 levels of the architecture 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 [Clark88] should not still stand. However, it is recognised that the end-to-end design principle is hedged around with stronger caveats than before. Also there have been some attempts at defining additional principles, and better articulating the concerns around the end-to-end design principle (see for instance [Ahlgren04, Ford09] and Clark's unpublished ideas on defining end-to-end in terms of trust-to-trust). But we must be open for understanding how we can establish architectural 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.

RESEARCH STYLE

The basis of scientific research lies in the ability of researchers 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 and, although some believe that a future Internet will result from agreement by committees representing industry players and governments, we do not believe that to be a realistic view. The present Internet overtook the plans of committees of experts 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 data networking research community and the Internet Engineering Task Force (IETF), but this is unlikely to be realistic either, simply because the nature and importance of the Internet have 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 and fora (including newly constituted ones). 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 an increasingly 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 both:

- **Those with a rationalist 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 exactly what is needed to develop completely new paradigms that address 'Missed Opportunities', but it is much less appropriate for solving the 'Immediate Problems' or for working on the 'Evolutionary Mechanisms'.

- **Those with an empiricist viewpoint:** 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' challenge. It requires rigorous experimental method, sharing of experimental 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 preferably multi-disciplinary, research teams, and long-lived institutions to foster consensus and provide continuity. And last but not least, it not only requires existing stake-holders to be consulted, those with commercial interests and those with operational experience, but also to enable potential future stakeholders having a voice in the overall 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 cultural perspectives, in national and international law, in the economics of deployment, in human rights, etc., not only experts in inter-domain policy-routing, radio fading 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 in the Internet at design-time.

TECHNOLOGY: LOOKING FORWARD AND LOOKING BACKWARD

The success of today's Internet is intrinsically coupled with the advancement of computing and communication technologies. To speculate where the future of Internet would lead us, although we cannot see clearly into the crystal ball, we can look back to see where we have come from, and learn a great deal from this exercise.

A Few Observations about the Internet's Success

Everyone agrees that Internet has been a tremendous success. However not everyone has thought about why; very often this success is attributed to the original design of the Internet architecture. Understanding this can be an important step towards answering the question of how the Future Internet should look.

In looking back, the first observation one can make is that *the success of the Internet is not only the success of individual Internet applications* (hardly any of today's applications except email was envisioned at the beginning) *but also the variety of developed application types*. The Internet has nurtured fantastic innovations that contributed to the welfare of the global society. Ever since the first killer app, World Wide Web, showed up in early 90's, the rate of innovative, successful applications has been accelerating over the last 15 years: eBay, Amazon, MySpace, Facebook, Skype, Google, BitTorrent, Webcams, RFID supply chain systems, mobile mapping applications, geotagging applications, ...

These successful applications all share one fact: none of them is the product of a small number of visionaries; rather their inventors were no-names till their inventions took over the world in no time. What will be the next killer app? We cannot hope to predict a priori, but we know for sure that it will come soon. Internet innovations are driven by the global user community, people who use the Internet everyday keep inventing new ways of using it. In the past, a number of large companies have invested into searching for "next killer app's", with little return. Thus one can make the second observation that *successful Internet inventions are by practitioners*. They understand their need, and they invent a solution to best suit the need.

What are the fundamental enablers that have unlatched the user community's creativity? This question leads to our third observation: *the last 15 years have handed the global user community two enablers: (1) the abundant and affordable computing resources provided by Moore's Law, and (2) universal connectivity provided by the Internet.* Not only the mere co-existence of these enablers has fuelled the Internet but its mutual re-enforcement. It is undeniable that the original Internet architecture was a successful design. However it is Moore's Law and cost reduction of storage that brought unprecedented success to the Internet by providing end users abundant and affordable computing resources, which led to the user innovations in Internet applications.

Are We Doing Science, or Engineering?

Networking Research as an Engineering Undertaking

As stated by Petroski, "Engineering has at its principal object not the given world but the world that engineers themselves create". Physicists study a given world to find out how it works. In contrast, network researchers started by first building networks. Networking research is all about what is the best design and how to build best networks, for whatever the definition of "best". But what is "best"? And how can we engineer systems that are "best" for a metric at some point without requiring a full re-engineering for another set of metrics of what "best" means for changed set of stakeholders.

Networking Research as a Scientific Endeavour

How do we find out the best network design? We can look back into the early days of networking to look 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 ARPAnet/NCP design and operation experience; the End-to-End principle became a foundation block in Internet architecture. The hour glass model also 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 overtime. Congestion control was not in the original protocol design, but was quickly engineered in when the congestion crisis arose in practice.

Internet is an artefact that can and should be studied just as physicists study the given world. The fact that we designed a network does not lead to the conclusion that we 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 protocols (in fact, a large number of them) exist in the system and we do not know how to derive the combined behaviour out of the interactions among various pieces.

Thus after the engineering phase (or even parallel with the engineering design), network researchers have to take on the role of a scientist to understand the world we just built. We need to rigorously follow a set of well established scientific research 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. But it also includes the ability to properly observe the phenomenon, which occur in daily operation. For that, an inherent ability is required to measure relevant 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 future designs of new artefacts or tunings of existing ones.

To summarize, we believe that the network research community must collectively play a double role in advancing the technology. 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, as we explain next.

Looking Back and Looking Forward: A Few Lessons

We need a full understanding of the driving forces behind the Internet's unprecedented success. The Internet would not have succeeded so greatly without Moore's Law. Computing technologies are moving forward with accelerated speed. The successful 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. How do we do it? 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, 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 in general 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 can and *should* change over time. All new systems start small. Once successful, they grow larger. The growth will bring the system to an entirely new environment that the original designers may not have envisioned, together with a new set of 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 its limitations (designed when the system was small, hence impossible to foresee all future threats). 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.

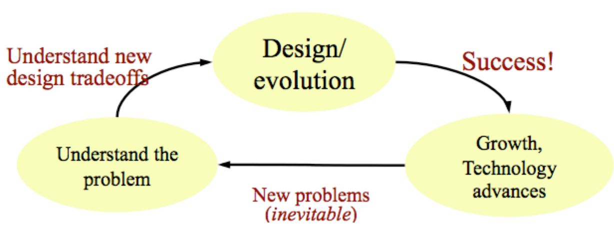

Last, but not the least, we need to develop a good understanding of networking architecture advances. It is often stated that the Internet architecture has not been changed over the last XX years (although depending how one defines *changes*, this statement may or may not be entirely true in reality [Handley06]). But how should changes be made? Would they arrive as a revolution, or should the adjustments being made continuously as in an evolutionary process, as the above figure indicated? In any case, the Internet architecture should be prepared to evolve either based on smooth evolution or based on rapid introduction of fundamental paradigms.

ECONOMICS: DESIGN FOR SUSTAINABILITY OR SUSTAININIG OUR CURRENT DESIGN?

The Internet was designed as a scalable packet routing fabric with a design 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 very edge of the network, with services like Google, MySpace, Facebook and many others being the drivers of this innovation. But this focus on edge innovation has left many problems unsolved, such as QoS, multicast and many others - problems that require economic coordination among many stakeholders. Consideration of deployability and economic sustainability for all market players is crucial. This however 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 this sentiment is similar in the economics community.

Hence, little work has been sprouting from the intersection of technology and economics, although this is changing slowly. However, there is a danger that this joint consideration of many of the grand challenges ahead is happening too slowly. Examples for these changes are the development of a resource accountability framework, the questions on scalability and dynamics and even fundamentals around security. It was identified in the Frankfurt think tank meeting that careful consideration needs to be given to determine the right audience for the debate that is needed. This goes beyond the problem of merely identifying the right people to invite to the debate. 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 the common debate.

But in addition to addressing the right community, we need to be aware that the large majority of the world population does not yet have access or is not actively using the Internet. Thus, there is still a significant untapped socio-economic potential of billions of consumers, entrepreneurs, innovators, creators, communicators, activists but also disruptors and criminals. The process of bringing the vast majority of the world population online will thus have significant impacts on both the online and real worlds, on technology as well as economic foundations of existing and future markets. For instance, research has shown the economic enabling power that mobile phones have in the developing world. A correlation between poverty reduction and the take up of mobile telephony is undisputed. Connectivity and Internet access is believed to generate similar positive effects, creating bigger markets, providing access to more information and new customers.

(Some) Grand Challenges That Will Require a Joint Debate

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., [EIFFEL06]. 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 do 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, as being pointed out in, e.g., [Gillett05], for the broadband industry. In order to sustain such healthy value chains, it is crucial to be able to account for the usage of resources along the value chain, 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 co-joint system design challenge.

Current Business Structures and Their Validity for the Future

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 but in particular any fundamental approach to change the current Internet will undoubtedly have an impact on these existing business 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 *here* to *there* are crucial when targeting a wide adoption of proposed changes.

A Way Forward

We need to address the grand challenges in economics. This needs to start with gathering the right audience for this work and it needs to be driven by a clear emphasis on the concrete problems and the quest for some answers to these. The abovementioned challenges give a first glimpse onto such problems. For the think tank, concrete topical work with concrete outcomes is seen as the best way forward. For this, themed sessions or interim working groups can be used in the near future.

SOCIETY: THE DUAL ROLES OF THE FUTURE INTERNET AND SOCIAL ORGANIZATION

As mentioned earlier, the Internet is 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 (reasonably immediate, but not requiring real-time end-to-end connectivity) 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, even attacking the infrastructure itself. 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 in society.

The demands of improving social communication and reflecting social structure, while at the same time increasingly addressing issues of privacy and safety in a completely 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. This process is repeated all the time. 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.

Addressing Information Overload

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?

Does Governance Have 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 that will have only a 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 that can transmit photos are having profound effects on the capability of individuals to constrain their governments at times when the governments may not want that. Even more than that, there are 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.

NEXT STEPS

This report is only the start. It is intended to stimulate, even provoke discussion on the major points of why and how we will be going about the Future Internet.

Along this line, the major next steps will be:

- Continue the discussion on Grand Challenges through
    - Deepening the topical discussions towards what is broken, identifying as to what needs fixing?
    - Identifying what we want preserve as fundamentals and what needs to be questioned?
    - Debate a design process that can serve as an evolutionary kernel for future designs
- Invite discussion on the challenges and their potential solutions
- Clearly formulate the outcome of these discussions in a form that can be disseminated to various stakeholders, most importantly to the research community itself.

The next steps outlined above are obviously aligned with the objectives of the EIFFEL support action. It was recognized in the first EIFFEL think tank meeting that a tool such as EIFFEL, gathering renowned experts and steering towards a clear formulation of consent and dissent on key topics of the Future Internet, is required. Hence, the most important next step can be seen as strengthening the idea of an independent (EIFFEL) think tank. This not only requires the facilitation by the EIFFEL caretakers but the involvement and contribution of the key experts being involved in this exercise.

REFERENCES

[Ahlgren04] B. Ahlgren, M. Brunner, L. Eggert, R. Hancock, S. Schmid, "Invariants: A New Design Methodology for Network Architectures" In: SIGCOMM 2004 Workshop on Future Directions in Network Architecture (FDNA 04), 30 Aug 2004, Portland, Oregon, USA.

[Clark88] D. D. Clark, "The design philosophy of the DARPA internet protocols," In: Proc. ACM SIGCOMM'88, Computer Communication Review 18 (4) pp. 106--114 (August, 1988).

[Clark05] D. Clark, K. Sollins, J. Wroclawski, R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet," In: IEEE/ACM Transactions on Networking 13 (3) pp. 462--475 (June, 2005).

[Dovrolis08] C. Dovrolis, "What would Darwin think about clean-slate architectures?" In: ACM SIGCOMM Computer Communication Review 38 (1) pp. 29--34 (2008).

[Ford09] A. Ford, P. Eardley, B. van Schewick, "New Design Principles for the Internet," In Int'l Workshop on the Network of the Future (Future-Net'09) (June 2009) (to appear).

[Tro05] D. Trossen (ed), "The Core-Edge Story", whitepaper of the Core-Edge Dynamics working group at the Communications Futures Program, MIT, 2005

[Hol08] R. Hollingsworth, K. Müller, "Transforming socio-economics with a new epistemology," Socio-Economic Review, vol. 6, pp. 395–426, 2008

[EIFFEL06] P. Mahonen (ed), D. Trossen, D. Papadimitrou, G. Polyzos, D. Kennedy, "Future Networked Society", EIFFEL whitepaper, 2006

[Gillett05] S. Gillett (ed), "The Broadband Incentive Problem", whitepaper of the Broadband working group at the Communications Futures Program, MIT, 2005

[Handley06] M. Handley, "Why the Internet Only Just Works," In BT Technology Journal Vol 24, No 3, July 2006

