



HAL
open science

Toward a polymorphic future internet: a networking science approach

Kavé Salamatian

► **To cite this version:**

Kavé Salamatian. Toward a polymorphic future internet: a networking science approach. IEEE Communications Magazine, 2011, 49 (10), pp.174 -178. 10.1109/MCOM.2011.6035832 . hal-00737885

HAL Id: hal-00737885

<https://hal.science/hal-00737885>

Submitted on 21 Nov 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Toward a Polymorphic Future Internet: A Networking Science Approach

Kavé Salamatian, Université de Savoie

ABSTRACT

In this article, I develop two major propositions. First, the future Internet should be polymorphic, and it should conciliate different architectural networking paradigms. The second proposition is that the future Internet should necessarily build on a strong theoretical basis from a networking science that is in the course of development. Particularly, I describe how virtualization makes possible a polymorphic future Internet and provides for easy deployment of new cooperation schemes. Then I attempt to expound on the aspects of security and scalability of the future Internet. This analysis also hopes to account for some justification as to the need for a clean-slate approach to fundamentally tackle the shortcomings of the current Internet and why a polymorphic future Internet is needed to for the coexistence of evolutionary and revolutionary schemes. Admittedly, the future Internet has to be built on strong foundations of networking science. A global theory of cooperation over networks is therefore most needed to foresee the important issues in the design and deployment of the future Internet. Methodologies for experimenting and validating large-scale mechanisms by reproducible large-scale experiments are, moreover, among the expected outcomes of the emerging networking science, which I also discuss in this article.

INTRODUCTION

Over the preceding 40 years, the Internet has grown to a network that connects an estimated 1.8 billion users that now has reached a penetration rate of almost 25 percent globally. What was once a tool known and used by only a small number of high-profile researchers has, within one generation, become a universal commodity, like electricity, in the daily lives of hundreds of millions of users within one generation. Such an unmatched phenomenon rightly justifies the paramount importance appended to the ongoing discussion on the “future Internet.” Academia and researchers around the globe now attempt to envision the definition, the design, and the construction of the future Internet and its imperatives. On the research side, for instance, among the different initiatives undertaken on the subject of “future Internet” worldwide, reference

could be made to the U.S.-based Global Environment for Network Innovation (GENI) initiative [1] and the European Union Future Internet Research and Experimentation (FIRE) initiative [2]. The United Nations World Summits on Information Society, held in 2007 and 2009, while addressing the theme of the Internet, further augmented its scope as the main component of the Information Society of the future and formally introduced cultural aspects into it as well. This flurry of interest in the future of the Internet is also a source of confusion; different and competitive requirements as well as architectural concepts are blurring our vision of it.

I give here my personal views, presented initially to the Kaleidoscope 2010 conference, and share them with the readership of *IEEE Communications Magazine* for further discussion and criticism. To trigger the discussion on my perspectives regarding the future Internet, I ask for the liberty to be a little more radical than usual, as I really think that it is important to be provocative, at least in some ways.

I have two positions in this article: first, that the future Internet should be polymorphic; that is, it will need to conciliate different architectural paradigms. Particularly, the future Internet (at least in a long transitory period) would have to support the evolution of the current Internet in the form of IPv6 or any other evolution of the current IP architecture along with other more revolutionary paradigms. Therefore, it is imperative to consider flexibility as the main property of the future Internet to provide smooth coexistence of evolutionary and revolutionary paradigms. My second position is that we should build the future Internet on the strong scientific foundations of networking science. The assumption of this science is that ancient networks, like roads, the postal service, and telephony, have some common principles with the latest developments of networks, like the Internet and online social networks. Networking science borrows some of its principles from well established disciplines, such as computer science, physics, social science, and economics, but it also maintains its own particular fundamental laws and principles. The multidisciplinary concept of networking science is therefore imperative to better understand and analyze the future Internet. Admittedly, networking science is under development, and ongoing research is still underway

to define all its dimensions. Yet there is an explicit consensus in academia on some basic concepts. Among the founding concepts of this new discipline, I wish to expound on the principle of cooperation as a reading lens for understanding the main issues of the future Internet.

COOPERATION: THE VERY ESSENCE OF NETWORKING

A network can be defined as a set of nodes that are cooperating with each other to distribute (exchange) information. A network architecture is there to provide a framework for nurturing this cooperation. The open systems interconnection (OSI) architecture allows only cooperation between layers in the same level of the architecture, and uses primitives provided by lower layers. This makes the development of new cooperation tractable and easy to formalize. Layer opacity and independence have enabled programmers to concentrate on a single layer and implement services without being obliged to tussle with other layers. However, layering always charges a performance cost. Indeed, one can imagine other types of cooperation, and the past decade has seen the development of such alternatives. For example, cross-layering (i.e., enabling a layer to access information and to interact with any other layers) for higher efficiency, performance, resource management, and security. In particular, we have witnessed the emergence of the *autonomic network* idea, whereby a network component is seen as an active element that is *self-conscious* and interacts with its environment. The development of the autonomic network concept has led to a major shift of the networking paradigm from a layered to a puzzle view, where autonomic components cooperate with each other.

The classical view tends to consider the role of a network element in processing packets only. By considering that the network element can decide which type of cooperation it will implement, networking science broadens its scope considerably. The decision is based not only on network objectives defined by the network operator or the service provider, but also on selfish goals of the person or entity that owns the node.¹ Selfishness might drive a node to go into standby mode to reduce battery consumption or drop packets that are not of interest. The observation that nodes might be selfish is a major change. It is motivated by several application scenarios, like ad hoc wireless networks, inter-AS routing, and peer-to-peer networks. Selfishness means a node cannot assume that its neighbors will fully cooperate, and it needs to react accordingly. Therefore, the node might have to implement a palette of different profiles for different contexts and appear polymorphic. Integrating selfishness is then a central motivation for proposing that the future Internet architecture be polymorphic.

As economics purports to address the production, distribution, and consumption of services and goods, networking science addresses the production, distribution, and consumption of “information.” Understanding, formalizing, and shaping cooperation for the distribution of infor-

mation are the major goals of network science. However, economics is a science that applies to resources that are scarce and have alternative uses. What differentiates digital information from other goods is that information is universal and infinitely reusable; that is, once produced, information supply becomes infinite, and there is almost no cost in reproducing it. Moreover, information is ambiguous (discriminating wrong from correct information is hard), whereas other goods and services are unambiguous. Therefore, a new cooperation theory, different from classical economical theory, must be developed to integrate network cooperation and the economical value of information.

TO CLEAN THE SLATE OR NOT? IS IT REALLY AN ISSUE?

The Internet was always designed, developed, and deployed simultaneously. More precisely, whenever a problem arose, a solution was proposed and analyzed, following a technical consensus at the Internet Engineering Task Force (IETF), resulting in a Request for Comments (RFC) that solved the issue or implemented a new feature. The IETF consensus guaranteed that the proposed evolution of the Internet was backward-compatible and complied with the sacrosanct axiom of “no harm to what works.” This process is one of the main explanations for the huge success and relative stability of the current Internet.

The clean-slate approach comes from the belief that it is impossible to resolve the challenges facing today's Internet without rethinking the fundamental assumptions and design decisions underlying its current architecture. The incremental approach changes the Internet architecture by backward-compatible patches; the clean-slate approach advocates out-of-the-box thinking with an architectural redesign, based on ameliorated concepts and abstractions to answer the current challenges.

Deciding between a clean slate and an evolutionary approach has become a major subject of controversy in the Internet community. Clean slate critics advocate that it is not wise to destruct something that works. Clean slate adepts state that the current Internet architecture constrains future evolution and hinders the deployment of radical solutions that attack existing problems (to be described later).

Indeed, the future Internet (at least in a long transitory period) should have to support the evolution of the current Internet. However, the future architecture should not hinder the design and deployment of other more revolutionary paradigms. This means that the key property of the future Internet architecture is having enough flexibility to enable the coexistence of the evolution of the current Internet (via incremental patches) with clean-slate revolutionary approaches. We should design the future Internet so that the question of cleaning the slate becomes irrelevant. I describe later why I believe that such flexibility is achievable with current virtualization technologies.

Nonetheless, imagining a new architecture is

Information is ambiguous (discriminating wrong from correct information is hard), whereas other goods and services are unambiguous. Therefore, a new cooperation theory, different from classical economical theory, must be developed to integrate network cooperation and the economical value of information.

¹ The owner of a computer in a peer-to-peer network or the owner of an autonomous system (AS) border gateway router.

There is a consensus in the research community and in the larger audience that the current Internet has some shortcomings that make its evolution and/or revolution inevitable. I give some of the main rationale along three directions here: flexibility, security, and scalability.

a tough task. We need to choose, among the large set of possible architectures, that could replace the current Internet. Going for a clean-slate approach means that it will not be possible to deploy it easily on the current Internet. This means that testing, validating, and debugging clean-slate approaches need large-scale experimental facilities. This explains why almost all initiatives on the future Internet, like [1, 3], are backed by large-scale platforms. However, the networking research community has not yet agreed on how to conduct reproducible large-scale experiments. Answering this question is one of the major challenges of network science and a prerequisite for validating clean-slate approaches. This issue is another one of the major motivations of why the future Internet needs to be polymorphic. As clean-slate approaches can be difficult to validate, one has to ensure a fallback of network nodes to classical IP technologies. Thus, at least for reliability reasons, one has anyway to consider the network as polymorphic.

FUTURE INTERNET MOTIVATIONS AND RATIONALES

There is broad consensus that the current Internet has some shortcomings that make its evolution and/or revolution inevitable. I give some of the main rationale along three directions here: flexibility, security, and scalability. This critical analysis sheds some light on the directions to go to design the future Internet.

FLEXIBILITY OR THE FUTURE INTERNET CONTORTIONIST

I have already presented some arguments as to why the future Internet has to be flexible enough to accommodate different cooperation models simultaneously. Another rationale relates to new application deployment. The current Internet does not provide architectural hooks to deploy services beyond the socket interface; developers have no access to routing and addressing. For this reason, during the past decade, developers have frequently raised routing and addressing into the application level, where they can access it. Peer-to-peer and overlay networks are examples of this approach, and implement a complete cooperation scheme in the application level (more precisely, above the socket interface). While this approach has been very successful, it is not optimal, as the packets still have to go through the underlying narrow hourglass of IP that acts as a bottleneck. A network where one could implement and deploy new network protocols or cooperation schemes without disturbing other running protocols would solve the application deployment issue and moreover provide an efficient way for innovative service deployment.

The network research community has pursued the quest for a flexible experimentation platform for future Internet research. This has resulted in the development of Planetlab [3]; its European counterpart, OneLab [4]; and GENI [1]. The flexibility in these experimental platforms was attained, thanks to the wide generalization of vir-

tualization approaches [5]. Virtualization ensures full isolation (fault, software, and performance isolation) between virtual machines. Because it enables the parallel execution of different networking systems (routing, addressing, etc.), it opens the way for a polymorphic future Internet.

Virtualization techniques encapsulate a full virtual machine into a single file that can be easily migrated and executed over any virtualized hardware. The encapsulation property opens the perspective of easy deployment of services just by distributing the encapsulated virtual machine. Thus, it is possible to implement the service over a large infrastructure of virtualized servers/routers. Last but not least, virtualization ensures interposition, because all activities go through a monitoring layer. Nowadays, commodity hardware has enough processing power to build virtual routers over clusters of multicore routers [6]. This opens the perspective of building realistic routers, implementing the polymorphic future Internet.

SECURITY: THE ACHILLES'S HEEL OF THE CURRENT INTERNET

One of the major rationales for the development of a future Internet is security. Indeed, the current Internet is plagued with spam, phishing, denial of service attacks, exploiting security breaches, and so on. However, only a small proportion of security-related problems could be traced back to the Internet architecture. Network security can be defined along three directions: communication security (ensuring that communication remains secret), cooperation security (ensuring that cooperation in networks is possible and secure), and, finally, application security (ensuring that an application is doing what it is supposed to do). Future Internet architecture should support each of the above components.

Security was not considered to be an essential component of the current Internet architecture. While support for communication security through IPSec exists (as an optional element in IPv4 and a mandatory one in IPv6), no support for the two other aspects of security exists. This has resulted in the current security status, where several external services, like virtual private networks (VPNs), firewalls, proxies, and Secure Sockets Layer (SSL), compensate for the lack of support built into the architecture.

The current status of Internet security has motivated extreme positions. On one hand, some suggest the integration of all security primitives inside the architecture, so that applications can fully rely on architectural security services. However, this could degrade the performance of network elements significantly. An opposite viewpoint advocates a "keep it simple, stupid" (KISS) approach, where the architecture should remain simple by focusing only on packet forwarding, and security mechanisms should remain on the network edges, like firewalls and IPSec mechanisms. This viewpoint is also problematic, as one cannot control the network completely without acting on the core of the architecture. In between these two extremes, the future Internet will have to determine the least common denominator of security support; in other terms, what should be mandatorily integrated into the archi-

ture and what could be considered as an applicative service that will cooperate with other components through the architecture.

The experience of IPsec shows that communication security can be supported satisfactorily through solutions that do not need to go deep into the core of the architecture. By looking at security through the cooperation lens, the future Internet architecture needs to integrate mechanisms to secure cooperation. This means that at least three basic security mechanisms have to be integrated to the future Internet architecture: an isolation mechanism shielding components running in the same execution environment strictly and at the deepest level, an authentication mechanism authenticating the source of a running code, and a monitoring mechanism that will evaluate the cooperation behavior of executing components and compare it with normal or expected behaviors. None of these mechanisms exists in the current Internet. However, virtualization provides mechanisms to guarantee fault, performance, and execution isolation. It also guarantees interposition. This means that a virtualized and polymorphic architecture for the future Internet will provide two out of the three needed basic security mechanisms. The scope of the authentication service is still a research topic, as it is not yet clear if a global authentication and/or identity mechanism is mandatory, or if only a local and trust-based scheme will be enough to cover the large spectrum of scenarios.

Permanent monitoring is a trade-off between performance and security; the more security we need, the more exhaustive the monitoring should be. We also have a trade-off between flexibility (in terms of the range of acceptable behaviors) and security; detecting abnormal behaviors entails reducing the range of acceptable behaviors to be able to differentiate them from abnormal ones.

It is worth recalling that security is a negative concept: you do not appreciate it when you have it; you only realize its significant importance when you have lost it. This means that rather than speaking about providing security, one should talk about reducing vulnerabilities. Almost 30 years of experience in Internet security has taught us that it is too costly and even impossible to remove all risks. The consequence of this is that we have to increase the resilience of the future Internet architecture to ensure survivability and to reduce the impact of security risks. In other terms, security risks should be assumed as a plausible operational hypothesis in the design of networked systems, and architectural solutions should be provided to detect and to contain them. This is a radically different position from the current approaches, where the emphasis is put on authentication of users through passwords/biometrics and assuming that authenticated users are entitled to do whatever they do. In the collaborative approach, we have to assume that users (even authenticated) can misbehave, and we should be able to detect and contain them. This is another shift in the security paradigm.

A noteworthy action that is supporting the paradigm change mentioned above is the ITU-Cybex action, aiming to develop a standard cybersecurity information exchange framework [7]. ITU-Cybex is based on the observation that

security risks are now internationally endemic. In this context, information about security events should be shared rather than being considered as private issues.

SCALABILITY, OR A DELUSION OF GRANDEUR

Scalability is another cardinal issue for the future Internet. The past decade has seen almost a doubling of Internet traffic per year. The size of routing tables, which are the main indicator of the complexity of the routing operation, has seen yearly growth of 19.4 percent from 2002 to 2008.²

Before dealing with scalability, we have to answer the question of why we need an address and how to answer this need. Initially, IP addresses were designed to identify resources on the network and locate them. However, this mixing of roles leads to some contradictions. On one hand, using an IP address as an identifier means that a unique IP address should be assigned to each resource in the network. On the other hand, using the IP address as a locator means that the address should change when the location of the resource changes. The locator address should satisfy some topological constraints (the locator address of nearby resources should be close), whereas identifiers should meet some semantic constraints (the addresses of all resources of the same operator should be close). The contradiction between the requirement of these two roles has resulted in the fragmentation of the IP address space and the scalability problem. While IPv6 solved the issue of identification address space exhaustion, it does not resolve the address space fragmentation that is the source of the scalability problem.

During the past two decades, several propositions have reduced address space fragmentation. In 1989, BGP introduced AS-level IP address aggregation; in 1993 classless Internet domain routing (CIDR) was used to regain the address waste resulting from past class-based allocations. However, address allocations had to remain compatible with previously allocated addresses because of the identification role of IP addresses. The Routing Research Group (RRG) of the Internet Research Task Force (IRTF) is currently investigating ways to improve the scalability of the current Internet in the current IP architectural framework. I do not give an exhaustive view of the proposed schemes. However, the main rationale of almost all of these proposals [8, 9] is to decouple the identifier and locator roles of the address. For example, Locator/Identifier Separation Protocol (LISP) [8] proposes a mapping system between the identifier and locator through an Alternative Topology (ALT) service. This mapping enables one to find the location of a given identifier, similar to what a Domain Name Service (DNS) server does. Identifier-Locator Network Protocol (ILNP) [9] has a similar mapping, but is done through DNSSEC rather than a specific service.

A “clean slate” approach to scalability does not enforce compatibility with the current IP architecture. New network scenarios, like delay-tolerant networks, showed that routing might not be possible in some scenarios. It was even shown that in order to optimize the throughput, network coding, which is not based on routing,

Scalability is another cardinal issue to consider for the future Internet. The past decade has seen almost a doubling of Internet traffic per year. The size of routing tables, which are the main indicator of the complexity of the routing operation, has seen yearly growth of 19.4 percent from 2002 to 2008.

² See BGP table size evolution at <http://bgp.potaroo.net/cidr/>.

Nearly a decade after most of the IPv6 standard was completed, the vast majority of software and hardware still uses IPv4. A polymorphic future Internet should provide for the co-existence of traditional IP addresses with more revolutionary addressing scheme and enable a gradual introduction of new routing schemes.

should be used. Fundamentally, addressing is a topological embedding adapted to a particular cooperation need; that is, it is a function that accepts an identifier and returns the position of the needed information to a topological space. It was shown that when the embedding is compact (i.e., when two nearby items are mapped by the addressing embedding into close addresses), addressing implies routing and vice versa. In other words, if one knows the identifier of what s/he wants, s/he can derive it from the path directly to reach it. This property means that scalable routing is possible and even trivial when compact embedding exists. IP (v4 or v6) addressing is not compact, as close nodes are not necessarily close in the IP address space. Nonetheless, compact embedding does exist. For example, the Content Addressable Network [10] (with the assumption of no node withdrawal) defines compact embedding. The question of knowing whether we can embed the particular addressing need of a specific cooperation scheme into compact embedding is one of the major questions of networking science. Peer-to-peer (P2P) and overlay networks have demonstrated that by lifting the IP addressing backward compatibility constraints, scalability can be achieved, and address fragmentation avoided. This advocates a clean-slate approach for the future Internet to enable deployment of new addressing/ routing schemes. Indeed, one can note that IP addresses are still needed, even on P2P and overlay networks. However, this is more a kind of link layer connectivity issue than a fundamental need of IP addressing.

A polymorphic future Internet should provide for the coexistence of traditional IP addresses with a more revolutionary addressing scheme and enable gradual introduction of new routing schemes. IPv6 showed the difficulty of introducing radical changes into the network. Nearly a decade after most of the IPv6 standard was completed, the vast majority of software and hardware still uses IPv4.

CONCLUSION

In this article, I define two main positions: the future Internet should be polymorphic, and it should be built on a strong networking science foundation.

I call for the development of a networking science that will provide the foundations for the future Internet. Networking science will be built mainly around the concept of cooperation. This cooperation theory is different from classical economics because it integrates network cooperation and the economic value of information. Another potentially high-impact area of networking science pertains to experimental validation. The networking research community still lacks largely agreed-on methodologies describing how to do reproducible large-scale experiments and how to validate large-scale mechanisms before deploying them in the wild. Answering these two needs is of the utmost importance for the future Internet.

I also recommended a polymorphic architecture for the future Internet. The rationale behind this proposal is that a polymorphic view reduces the importance of the “clean slate vs. evolutionary” controversy greatly. In particular, a polymor-

phic architecture will make the gradual deployment of noncompatible clean slate solutions possible along with IP-based evolutionary developments. Virtualization techniques can provide the flexible technology needed to build such a polymorphic future Internet. Moreover, these techniques will ensure isolation and interposition, two properties that are needed in a future architecture. The security paradigm would have to shift from schemes based only on encryption/authentication to monitoring schemes that permanently observe the cooperation behavior of network elements. The last topic developed in the article is scalability. My position is to promote infinitely scalable addressing and routing schemes by using suitable embeddings. This is an argument in favor of a clean-slate approach. Thus, the future Internet will support different addressing schemes suitable for different applications.

In summary, we should aim at designing a polymorphic Internet architecture that will let us benefit from innovations without harming the existing and successful current Internet.

ACKNOWLEDGMENTS

At the end, I would like to thank Serge Fdida for interesting discussions about the architecture of the future Internet, Dr. Mostafa Hashem Sherif for helpful comments, and the anonymous reviewers for their valuable advice.

REFERENCES

- [1] C. Elliott and A. Falk, “An Update on the GENI Project,” *SIGCOMM Comp. Commun. Rev.*, vol. 39, no. 3, 2009, pp. 28–34.
- [2] A. Gavras et al., “Future Internet Research and Experimentation: the FIRE Initiative,” *SIGCOMM Comp. Commun. Rev.*, vol. 37, July 2007, pp. 89–92.
- [3] L. Paterson and T. Roscoe, “The Design Principles of PlanetLab,” *Op. Sys. Rev.*, vol. 40, no. 1, Jan. 2006, pp. 11–16.
- [4] A. de la Oliva et al., “A Multihoming Architecture for OneLab,” *Proc. Real Overlays and Distrib. Sys. Wksp.*, Warsaw, Poland, July 2007.
- [5] T. Anderson et al., “Overcoming the Internet Impasse through Virtualization,” *IEEE Computer*, vol. 38, no. 4, 2005, pp. 34–41.
- [6] N. Egi et al., “Towards High Performance Virtual Routers on Commodity Hardware,” *CoNEXT*, 2008, p. 20.
- [7] A. Rutkowski et al., “CYBEX: the Cybersecurity Information Exchange Framework (x.1500),” *SIGCOMM Comp. Commun. Rev.*, vol. 40, Oct. 2010, pp. 59–64.
- [8] D. Farinacci et al., “Locator/ ID Separation Protocol (LISP),” Internet draft, Mar. 2009.
- [9] R. Atkinson, S. Bhatti, and S. Hailes, “Evolving the Internet Architecture Through Naming,” *IEEE JSAC*, vol. 28, Oct. 2010, pp. 1319–25.
- [10] S. Ratnasamy et al., “A Scalable Content-Addressable Network,” *SIGCOMM’01: Proc. 2001 Conf. Apps., Technologies, Architectures, and Protocols for Comp. Commun.*, New York, NY, 2001, ACM, pp. 161–72.

BIOGRAPHIES

KAVÉ SALAMATIAN (kave.salamatian@univ-savoie.fr) is a full professor in computer science at the University of Savoie. His main areas of researches are Internet measurement and modeling, and networking information theory. He was previously a reader at Lancaster University, United Kingdom, and associate professor at University Pierre et Marie Curie. He graduated in 1998 from Paris SUD-Orsay University where he worked on joint source channel coding applied to multimedia transmission over Internet for his Ph.D. In a former life, he graduated with an M.B.A., and worked on the market floor as a risk analyst and enjoyed being an urban traffic modeler for some years. He is working these days on figuring out if networking is a science or just a hobby; and, if it is a science, what are its fundamentals.