

Towards a Future Internet Architecture

Theodore Zahariadis¹, Dimitri Papadimitriou², Hannes Tschofenig³, Stephan Haller⁴,
Petros Daras⁵, George D. Stamoulis⁶, and Manfred Hauswirth⁷

¹Synelixis Solutions Ltd/TEI of Chalkida, Greece
zahariad@{synelixis.com, teihal.gr}

²Alcatel-Lucent, Belgium
dimitri.papadimitriou@alcatel-lucent.com

³Nokia Siemens Networks, Germany
hannes.tschofenig@nsn.com

⁴SAP, Germany
stephan.haller@sap.com

⁵Center of Research and Technology Hellas/ITI, Greece
daras@iti.gr

⁶Athens University of Economics and Business, Greece
gstamoul@aueb.gr

⁷Digital Enterprise Research Institute, Ireland
manfred.hauswirth@deri.org

Abstract. In the near future, the high volume of content together with new emerging and mission critical applications is expected to stress the Internet to such a degree that it will possibly not be able to respond adequately to its new role. This challenge has motivated many groups and research initiatives worldwide to search for structural modifications to the Internet architecture in order to be able to face the new requirements. This paper is based on the results of the Future Internet Architecture (FIArch) group organized and coordinated by the European Commission (EC) and aims to capture the group's view on the Future Internet Architecture issue.

Keywords: Internet Architecture, Limitations, Processing, Handling, Storage, Transmission, Control, Design Objectives, EC FIArch group.

1 Introduction

The Internet has evolved from a remote access to mainframe computers and slow communication channel among scientists to the most important medium for information exchange and the dominant communication environment for business relations and social interactions. Billions of people all over the world use the Internet for finding, accessing and exchanging information, enjoying multimedia communications, taking advantage of advanced software services, buying and selling, keeping in touch with family and friends, to name a few. The success of the Internet has created even higher hopes and expectations for new applications and services, which the current Internet may not be able to support to a sufficient level. It is expected that the number

of nodes (computers, terminals mobile devices, sensors, etc.) of the Internet will soon grow to more than 100 billion [1]. Reliability, availability, and interoperability required by new networked services, and this trend will escalate in the future. Therefore, the requirement of increased robustness, survivability, and collaborative properties is imposed to the Internet architecture. In parallel, the advances in video capturing and content/media generation have led to very large amounts of multimedia content and applications offering immersive experiences (e.g., 3D videos, interactive environments, network gaming, virtual worlds, etc.) compared to the quantity and type of data currently exchanged over the Internet. Based on [2], out of the 42 Exabytes (10^{18}) of consumer Internet traffic likely to be generated every month in 2014, 56% will be due to Internet video, while the average monthly consumer Internet traffic will be equivalent to 32 million people streaming Avatar in 3D, continuously, for the entire month.

All these applications create new demands and requirements, which to a certain extent can be addressed by means of “over-dimensioning” combined with the enhancement of certain Internet capabilities over time. While this can be a satisfactory (although sometimes temporary) solution in some cases, analyses have shown [3],[4] that increasing the bandwidth on the backbone network will not suffice due to new qualitative requirements concerning, for example, highly critical services such as e-health applications, clouds of services and clouds of sensors, new social network applications like collaborative 3D immersive environments, new commercial and transactional applications, new location-based services and so on.

In other words, the question is to determine if the architecture and its properties might become the limiting factor of Internet growth and of the deployment of new applications. For instance, as stated in [5] “*the end-to-end arguments are insufficiently compelling to outweigh other criteria for certain functions such as routing and congestion control*”. On the other hand, the evolution of the Internet architecture is carried out by means of incremental and reactive additions [6], rather than by major and proactive modifications. Moreover, studies on the impact of research results have shown that better performance or richer functionality implying an architectural change define necessary but not sufficient conditions for such change in the Internet architecture and/or its components. Indeed, the Internet architecture has shown since so far the capability to overcome such limits without requiring radical architectural transformation. Hence, before proposing or designing a new Internet Architecture (if a new one is needed), it is necessary to demonstrate the fundamental limits of the current architecture [7]. Thus, scientists and researchers from both the industry and academia worldwide are working towards understanding these architectural limits so as to progressively determine the principles that will drive the Future Internet architecture that will adequately meet at least the abovementioned challenges [EIFFEL], [4WARD], [COAST].

The Future Internet as a global and common communication and distributed information system may be considered from various interrelated perspectives: the networks and shared infrastructure perspective, the services and application perspective as well as the media and content perspective. Significant efforts world-wide have already been devoted to investigate some of its pillars [8] [9] [10] [11] [12] [13]. In

Europe, a significant part of the Information and Communication Technology (ICT) of the Framework Program 7 is devoted to the Future Internet [14]. Though many proposals for a Future Internet Architecture have already been developed, no specific methodology to evaluate the efficiency (and the need) for such architecture proposals exist. The purpose of this paper is to capture the view of the Future Internet Architecture (FIArch) group organized and coordinated by the European Commission.

Since so far, the FIArch group has identified and reached some understanding and agreement on the different types of limitations of the Internet and its architecture. Interested readers may also refer to [15] for more information¹.

2 Definitions

Before describing the approach followed by the FIArch Group, we define the terms used in our work. Based on [16], we define as “*architecture*” a set of functions, states, and objects/information together with their behavior, structure, composition, relationships and spatio-temporal distribution. The specification of the associated functional, object/ informational and state models leads to an architectural model comprising a set of components (i.e. procedures, data structures, state machines) and the characterization of their interactions (i.e. messages, calls, events, etc.).

We also qualify as a “*fundamental limitation*” of the Internet architecture a functional, structural, or performance restriction or constraint that cannot be effectively resolved with current or clearly foreseen “architectural paradigms” as far as our understanding/knowledge goes. On the other hand, we define as “*challenging limitation*” a functional, structural, or performance restriction or constraint that could be resolved as far as our understanding/knowledge goes by replacing and/or adding/removing a component of the architecture so that this would in turn change the global properties of the Internet architecture (e.g. separation of the locator and identifier role of IP addresses).

In the following, we use the term “*data*” to refer to any organized group of bits a.k.a. data packets, data traffic, information, content (audio, video, multimedia), etc. and the term “*service*” to refer to any action performed on data or other services and the related Application Programming Interface (API).² Note however that this document does not take position on the localization and distribution of these APIs.

3 Analysis Approach

Since its creation, the Internet is driven by a small set of fundamental design principles rather than a formal architecture that is created on a whiteboard by a standardization or research group. Moreover, the necessity for backwards compatibility and the trade-off between Internet redesign and proposing extensions, enhancements and re-engineering of today’s Internet protocols are heavily debated.

¹ Interested readers may also search for updated versions at the FIArch site:
http://ec.europa.eu/information_society/activities/foi/research/fiarch/index_en.htm

² The definition of service does not include the services offered by humans using the Internet

The emergence of new needs at both functional and performance levels, the cost and complexity of Internet growth, the existing and foreseen functional and performance limitations of the Internet's architectural principles and design model put the following elementary functionalities under pressure:

- **Processing/handling of “data”**: refers to forwarders (e.g. routers, switches, etc.), computers (e.g., terminals, servers, etc.), CPUs, etc. and handlers (software programs/routines) that generate and treat as well as query and access data.
- **Storage of “data”**: refers to memory, buffers, caches, disks, etc., and associated logical data structures.
- **Transmission of “data”**: refers to physical and logical transferring/exchange of data.
- **Control of processing, storage, transmission of systems and functions**: refers to the action of observation (input), analysis, and decision (output) whose execution affects the running conditions of these systems and functions. Note that by using these base functions, the data communication function can be defined as the combination of processing, storage, transmission and control functions applied to “data”. The term control is used here to refer to control functionality but also management functionality, e.g. systems, networks, services, etc.

For each of the above functionalities, the FIArch group has tried to identify and analyze the presumed problems and limitations of the Internet. This work was carried out by identifying an extensive list of limitations and potentially problematic issues or missing functionalities, and then selecting the ones that comply with the aforementioned definition of a fundamental limitation.

3.1 Processing and Handling Limitations

The fundamental limitations that have been identified in this category are:

- i. The Internet does not allow hosts to diagnose potential problems and the network offers little feedback for hosts to perform root cause discovery and analysis. In today's Internet, when a failure occurs it is often impossible for hosts to describe the failure (what happened?) and determine the cause of the failure (why it happened?), and which actions to take to actually correct it. The misbehavior that may be driven by pure malice or selfish interests is detrimental to the cooperation between Internet users and providers. Non-intrusive and non-discriminatory means to detect misbehavior and mitigate their effects while keeping open and broad accessibility to the Internet is a limitation that is crucial to overcome [16].
- ii. *Lack of data identity is damaging the utility of the communication system.* As a result, data, as an ‘economic object’, traverses the communication infrastructure multiple times, limiting its scaling, while lack of content ‘property rights’ (not only author- but also usage-rights) leads to the absence of a fair charging model.
- iii. Lack of methods for dependable, trustworthy processing and handling of network and systems infrastructure and essential services in many critical environments, such as healthcare, transportation, compliance with legal regulations, etc.

- iv. *Real-time processing.* Though this is not directly related to the Internet Architecture itself, the limited capability for processing data on a real-time basis poses limitations in terms of the applications that can be deployed over the Internet. On the other hand, many application areas (e.g. sensor networks) require real-time Internet processing at the edges nodes of the network.

3.2 Storage Limitations

The fundamental restrictions that have been identified in this category are:

- i. *Lack of context/content aware storage management:* Data are not inherently associated with knowledge of their context. This information may be available at the communication end-points (applications) but not when data are in transit. So, it is not feasible to make efficient storage decisions that guarantee fast storage management, fast data mining and retrieval, refreshing and removal optimized for different types of data [18].
- ii. *Lack of inherited user and data privacy:* In case data protection/ encryption methods are employed (even using asymmetric encryption and public key methods), data cannot be efficiently stored/handled. On the other hand, lack of encryption, violates the user and data privacy. More investigations into the larger privacy and data-protection ecosystem are required to overcome current limits of how current information systems deal with privacy and protection of information of users, and develop ways to better respect the needs and expectations [30], [31], [32]
- iii. *Lack of data integrity, reliability and trust,* targeting the security and protection of data; this issue covers both unintended disclosure and damage to integrity from defects or failures, and vulnerabilities to malicious attacks.
- iv. *Lack of efficient caching & mirroring:* There is no inherited method for on-path caching along the communication path and mirroring of content compared to off-path caching that is currently widely used (involving e.g. connection redirection). Such methods could deal with issues like flash crowding, as the onset of the phenomenon will still cause thousands of cache servers to request the same documents from the original site of publication.

3.3 Transmission Limitations

The fundamental restrictions that have been identified in this category are:

- i. *Lack of efficient transmission of content-oriented traffic:* Multimedia content-oriented traffic comprises much larger volumes of data as compared to any other information flow, while its inefficient handling results in retransmission of the same data multiple times. Content Delivery Networks (CDN) and more generally architectures using distributed caching alleviate the problem under certain conditions but can't extend to meet the Internet scale [19]. Transmission from centralized locations creates unnecessary overheads and can be far from optimal when massive amounts of data are exchanged.

- ii. *Lack of integration of devices with limited resources to the Internet as autonomous addressable entities.* Devices in environments such as sensor networks or even nano-networks/smart dust as well as in machine-to-machine (M2M) environments operate with such limited processing, storage and transmission capacity that only partly run the protocols necessary in order to be integrated in the Internet as *autonomous addressable entities*.
- iii. *Security requirements of the transmission links:* Communications privacy does not only mean protecting/encrypting the exchanged data but also not disclosing that communication took place. It is not sufficient to just protect/encrypt the data (including encryption of protocols/information/content, tamper-proof applications etc) *but also protect the communication* itself, including the relation/interaction between (business or private) parties.

3.4 Control Limitations

The fundamental limitations that have been identified in this category are:

- i. *Lack of flexibility and adaptive control*³⁴. In the current Internet model, design of IP (and more generally communication) control components have so far being driven exclusively by i) cost/performance ratio considerations and ii) pre-defined, static, and open loop control processes. The first limits the capacity of the system to adapt/react in a timely and cost-effective manner when internal or external events occur that affect its value delivery; this property is referred to as flexibility [20][21]. Moreover, the current trend in unstructured addition of ad-hoc functionality to partly mitigate this lack of flexibility has resulted in increased complexity and (operational and system) cost of the Internet. Further, to maintain/sustain or even increase its value delivery over time, the Internet will have to provide flexibility in its functional organization, adaptation, and distribution. Flexibility at run time is essential to cope with the increasing uncertainty (unattended and unexpected events) as well as breadth of expected events/ running conditions for which it has been initially designed. The latter results in such a complexity that leaves no possibility for individual systems to adapt their control decisions and tune their execution at running time by taking into account their internal state, its activity/behavior as well as the environment/external conditions.
- ii. *Improper segmentation of data and control.* The current Internet model segments (horizontally) data and control, whereas from its inception the control functionality has a transversal component. Thus, on one hand, the IP functionality isn't limited anymore to the "network layer", and on the other, IP is not totally decoupled from the underlying "layers" anymore (by the fact IP/MPLS and underlying layers

³ Some may claim that this limitation is "very important" or "very challenging" but not a "fundamental" one. As we consider it significant anyway, we include it here for the sake of completeness.

⁴ This limitation is often named by the potential approach aimed to address it, including autonomic networking, self-management, etc. However, none of them has shown ability to support flexibility at run time to cope with increasing uncertainty (since the control processes they accommodate are still those pre-determined at design time).

share the same control instance). Hence, the hour-glass model of the Internet does not account for this evolution of the control functionality when considered as part of the design model.

- iii. *Lack of reference architecture of the IP control plane.* The IP data plane is itself relatively simple but its associated control components are numerous and sometimes overlapping, as a result of the incremental addition of ad-hoc control components over time, and thus their interactions are becoming more and more complex. This leads to detrimental effects for the controlled entities, e.g., failures, instability, inconsistency between routing and forwarding (leading to e.g. loops) [22][23].
- iv. *Lack of efficient congestion control.* Congestion control cannot be realized as a pure end-to-end function: congestion is an inherent network phenomenon that can only be resolved efficiently by some cooperation of end-systems and the network, since it is a shared communication infrastructure. Hence, substantial benefit could be expected by further assistance from the network, but, on the other hand, such network support could lead to duplication of functions, which may harmfully interact with end-to-end principle and resulting protocol mechanisms. Addressing effectively the trade-off of network support without decreasing its scaling properties by requiring maintenance of per-flow state is one of the Internet's main challenges [16].

3.5 Limitations That May Fall in More than One Category

Certain fundamental limitations of current Internet may fall in more than one category. Examples of such limitations include:

- i. *Traffic growth vs heterogeneity in capacity distribution:* Hosts connected to the Internet do not have the possibility to enforce the path followed by their traffic. Hence, even if multiple alternatives to reach a given destination would be offered to the host, they are unable to enforce their decision across the network. On the other hand, as the Internet enables any-to-any connectivity, there is no effective means to predict the spatial distribution of the traffic within a timescale that would allow providers to install needed capacity when required or at least expected to prevent overload of certain network segments. This results into serious capacity shortage (and thus congestion) over certain segments of the network. Especially, the traffic exchange points (as well as certain international and the transatlantic links) are in many cases significantly overloaded. In some cases, building out more capacity to handle this new congestion may be infeasible or unwarranted. Two main types of limitations are seen in this respect: i) not known scalable means to overcome the result of network infrastructure abstraction, and ii) those related to congestion and diagnosability. These are related to at least the base functions of control and processing/handling.
- ii. *The current inter-domain routing system is reaching fundamental limits* in terms of routing table scalability but also adaptation to topology and policy dynamics (perform efficiently under dynamic network conditions) that in turn impact its convergence, and robustness/stability properties. Both dimensions increase memory requirements but also the processing capacity of routing engines [23][7] Related projects: [EULER] [ResumeNet].

- iii. *Scaling to deal with flash crowding.* The huge number of (mobile) terminals combined with a sudden peak in demand for a particular piece of data may result in phenomena that cannot be handled; such phenomena can be related at to all the base functions.
- iv. The amount of foreseen data and information⁵ requires significant *processing power / storage / bandwidth for indexing / crawling and (distributed) querying* and also solutions for *large scale / real-time data mining / social network analysis*, so as to achieve successful retrieval and integration of information from an extremely high number of sources across the network. All the aforementioned issues imply the need for addressing new architectural challenges capable to cope with the fast and scalable identification and discovery of and access to data. *The exponential growth of information makes it increasingly harder to identify relevant information (“drowning in information while starving for knowledge”). This information overload becomes more and more acute and existing search and recommendation tools are not filtering and ranking the information adequately and lack the required granularity (document-level vs. individual information item).*
- v. *Security of the whole Internet Architecture.* The Internet architecture is not intrinsically secure and is based on add-ons to, e.g. protocols, to secure itself. The consequence is that protocols may be secure but the overall architecture is not self-protected against malicious attacks.
- vi. *Support of mobility* when using IP address as both network and host identifier but also TCP connection identifier results in Transmission Control Protocol (TCP) connection continuity problem. Its resolution requires decoupling between the identifier of the position of the mobile host in the network graph (network address) from the identifier used for the purpose of TCP connection identification. Moreover, when mobility is enabled by wireless networks, packets can be dropped because of corruption loss (when the wireless link cannot be conditioned to properly control its error rate or due to transient wireless link interruption in areas of poor coverage), rendering the typical reaction of congestion control mechanism of TCP inappropriate. As a result, non-congestive loss may be more prevalent in these networks due to corruption loss. This limitation results from the existence of heterogeneous links, both wired and wireless, yielding a different trade-off between performance, efficiency and cost, and affecting several base functions again.

4 Design Objectives

The purpose of this section is to document the design objectives that should be met by the Internet architecture. We distinguish between “high-level” and “low-level” design objectives. High-level objectives refer to the cultural, ethical, socio-economic, but also technological expectations to be met by the Internet as global and common information communication system. High-level objectives are documented in [15]. By

⁵ Eric Schmidt, the CEO of Google, the world’s largest index of the Internet, estimated the size at around 5 million terabytes of data (2005). Eric commented that Google has indexed roughly 200 terabytes of that is 0,004% of the total size.

low-level design objectives, we mean here the functional and performance properties as well as the structural and quality properties that the architecture of this global and common information communication system is expected to meet. From the previous sections, some of low-level objectives are met and others are not by the (present) architecture of the Internet. We also emphasize here that these objectives are commonly shared by the Internet community at large

The remaining part of this Section translates a first analysis of the properties that should be met by the Internet architecture starting from the initial of objectives as enumerated in various references (see [27], [28], [29]). One of the key challenges is thus to determine the necessary addition/improvement of current architecture principles and the improvement (or even removal of architectural components needed to eliminate or at least tangibly mitigate/avoid the known effects of the fundamental limitations. It is to be emphasized that a great part of research activities in this domain consists in identifying hidden relationships and effects.

As explained in [27], the Internet architecture has been structured around eight foundational objectives: i) to connect existing networks, ii) survivability, iii) to support multiple types of services, iv) to accommodate a variety of physical networks, v) to allow distributed management, vi) to be cost effective, vii) to allow host attachment with a low level of effort and, viii) to allow resource accountability. Moreover, RFC 1287, published in 1991 by the IAB [36], underlines that the Internet architecture needs to be able to scale to 10^9 IP networks recognizing the need to add scalability as a design objective. In this context, the followed approach consists of starting from the existing Internet design objectives compared to the approach that would consist of applying a tabula rasa approach, i.e., completely redefine from scratch the entire set of Internet design objectives.

Based on previous sections, the present section describes the design objectives that are currently met, partly met or not met at all by the current architecture. In particular, the low-level design objectives of the architecture are to provide:

- **Accessibility** (open and by means of various/heterogeneous wireless/radio and wired interfaces) to the communication network but also to heterogeneous data, applications, and services, nomadicity, and mobility (while providing means to maintain continuity of application communication exchanges when needed). Accessibility and nomadicity are currently addressed by current Internet architecture. On the other hand, mobility is still realized in most cases by means of dedicated/separated architectural components instead of Mobile IP. see Subsection 3.5. Point 6
- **Accountability** of resource usage and security without impeding user privacy, utility and self-arbitration: see Subsection.3.1.Point.2
- **Manageability**, implying distributed, organic, automated, and autonomic/self-adaptive operations: see Subsection 3.5 and **Diagnosability** (i.e. root cause detection and analysis): see Subsection.3.1.Point.1
- **Transparency**, i.e. the terminal/host is only concerned with the end-to-end service; in the current Internet this service is the connectivity even if the notion of “service” is not embedded in the architectural model of the Internet: initially addressed but loosing ground.

- ***Distribution of processing, storage, and control functionality and autonomy*** (organic deployment): addressed by current architecture; concerning storage and processing, several architectural enhancements might be required, e.g. for the integration of distributed but heterogeneous data and processes.
- ***Scalability***, including routing and addressing system in terms of number of hosts/terminals, number of shared infrastructure nodes, etc. and management system: - see Subsection.3.5.Point.2
- ***Reliability***, referring here to the capacity of the Internet to perform in accordance to what it is expected to deliver to the end-user/hosts while coping with a growing number of users with increasing heterogeneity in applicative communication needs.
- ***Robustness/stability, resiliency, and survivability***: see Subsection.3.5.Point.2
- ***Security***: see Subsection.3.5 point 5, Subsection 3.1.Point.2 and 3.
- ***Generality*** e.g. support of plurality of applications and associated data traffic such as non/real-time streams, messages, etc., independently of the shared infrastructure partitioning/divisions, and independently of the host/terminal: addressed and to be reinforced (migration of mobile network to IPv6 Internet, IPTV moving to Internet TV, etc.) otherwise leading to segmentation and specialization per application/service.
- ***Flexibility***, i.e. capability to adapt/react in a timely and cost-effective manner upon occurrence of internal or external events that affect its value delivery, and ***Evolvability*** (of time variant components): not addressed see Subsection 3.4.Point.1.
- ***Simplicity and cost-effectiveness***: deeper analysis is needed but simplicity seems to be progressively decreasing see Section 3.4 Point 3. Note that simplicity is explicitly added as a design objective to -at least- prevent further deterioration of the complexity of current architecture (following the “Occam's razor principle”). Indeed, lowering complexity for the same level of performance and functionality at a given cost is a key objective.
- ***Ability to offer information-aware transmission and distribution***: Subsection 3.3, Point 1, and Subsection 3.5, Point 4.

5 Conclusions

In this article we have identified fundamental limitations of Internet architecture following a systematic investigation thereof from a variety of different viewpoints. Many of the identified fundamental limitations are not isolated but strongly dependent on each other. Increasing the bandwidth would significantly help to address or mitigate some of these problems, but would not solve their root cause. Other problems would nevertheless remain unaddressed. The **transmission** can be improved by utilizing better **data processing and handling** (e.g. network coding, data compression, intelligent routing) and better **data storage** (e.g. network/terminals caches, data centers/mirrors etc.), while the overall Internet performance would be significantly improved by **control and self-*** functions. As an overall finding we may conclude the following: **Extensions, enhancements and re-engineering of today's Internet protocols may solve several challenging limitations. Yet, addressing the fundamental limitations of the Internet architecture is a multi-dimensional and challenging research topic. While improvements are needed in each dimension, these should be combined by undertaking a holistic approach of the problem space.**

Acknowledgements. This article is based on the work that has been carried out by the EC Future Internet Architecture (FIArch) group (to which the authors belong), which is coordinated by the EC FP7 Coordination and Support Actions (CSA) projects in the area of Future Internet: NextMedia, IOT-I, SOFI, EFFECTS+, EIFFEL, Chorus+, SESERV and Paradiso 2, and supported by the EC Units D1: Future Networks, D2: Networked Media Systems, D3: Software & Service Architectures & Infrastructures, D4: Networked Enterprise & Radio Frequency Identification (RFID) and F5: Trust and Security. The authors would like to acknowledge and thank all members of the group for their significant input and the EC Scientific Officers Isidro Laso Ballesteros, Jacques Babot, Paulo De Sousa, Peter Friess, Mario Scillia, Arian Zweegers for coordinating the activities.

The authors would like also to acknowledge the FI architectural work performed under the project FP7 COAST ICT-248036 [COAST].

References

- [1] AKARI Project: New Generation Network Architecture AKARI Conceptual Design (ver1.1). AKARI Architecture Design Project, Original Publish (Japanese) June 2008, English Translation October 2008, Copyright © 2007-2008 NICT (2008)
- [2] Medeiros, F.: ICT 2010: Digitally Driven, Brussels, 29 September 2010, Source Cisco VNL (2010)
- [3] Mahonen, P. (ed.), Trossen, D., Papadimitrou, D., Polyzos, G., Kennedy, D.: Future Networked Society., EIFFEL whitepaper (Dec. 2006)
- [4] Jacobson, V., Smetters, D., Thornton, J., Plass, M., Briggs, N., Braynard, R.: Networking Named Content. Proceeding of ACM CoNEXT 2009. Rome, Italy (December 2009)
- [5] Moors, T.: A critical review of “End-to-end arguments in system design”. In: Proceedings of IEEE International Conference on Communications (ICC) 2002, New-York City (New Jersey), USA (April/May 2002)
- [6] RFC 1958: The Internet and its architecture have grown in evolutionary fashion from modest beginnings, rather than from a Grand Plan
- [7] Li, T. (ed.): Design Goals for Scalable Internet Routing. Work in progress, draft-irtf-rrg-design-goals-02 (Sep. 2010)
- [8] <http://www.nsf.gov/pubs/2010/nsf10528/nsf10528.htm>
- [9] http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=503325
- [10] <http://www.nets-find.net>
- [11] <http://www.geni.net/?p=1339>
- [12] <http://akari-project.nict.go.jp/eng/overview.htm>
- [13] http://mmlab.snu.ac.kr/fiw2007/presentations/architecture_tschoi.pdf
- [14] <http://www.future-internet.eu/>
- [15] FIArch Group: Fundamental Limitations of Current Internet and the path to Future Internet (December 2010)
- [16] Perry, D., Wolf, A.: Foundations for the Study of Software Architecture. ACM SIGSOFT Software Engineering Notes 17, 4 (1992)
- [17] Papadimitriou, D., et al. (eds.): Open Research Issues in Internet Congestion Control. Internet Research Task Force (IRTF), RFC 6077 (February 2011)
- [18] Akhlaghi, S., Kiani, A., Reza Ghanavati, M.: Cost-bandwidth tradeoff in distributed storage systems (published on-line). ACM Computer Communications 33(17), 2105–2115 (2010)
- [19] Freedman, M.: Experiences with CoralCDN: A Five-Year Operational View. In: Proc. 7th USENIX/ACM Symposium on Networked Systems Design and Implementation (NSDI '10) San Jose, CA (May 2010)

- [20] Dobson, S., et al.: A survey of autonomic communications. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* 1(2), 223–259 (2006)
- [21] Gelenbe, E.: Steps toward self-aware networks. *ACM Communications* 52(7), 66–75 (2009)
- [22] Evolving the Internet, Presentation to the OECD (March 2006), <http://www.cs.ucl.ac.uk/staff/m.handley/slides/>
- [23] Meyer, D., et al.: Report from the IAB Workshop on Routing and Addressing, IETF, RFC 4984 (Sep. 2007)
- [24] Mahonen, P. (ed.), Trossen, D., Papadimitrou, D., Polyzos, G., Kennedy, D.: *Future Networked Society*. EIFFEL whitepaper (Dec. 2006)
- [25] Trosse, D.: Invigorating the Future Internet Debate. *ACM SIGCOMM Computer Communication Review* 39(5) (2009)
- [26] Eggert, L.: Quality-of-Service: An End System Perspective. In: *MIT Communications Futures Program – Workshop on Internet Congestion Management, QoS, and Interconnection*, Cambridge, MA, USA, October 21-22 (2008)
- [27] Ratnasamy, S., Shenker, S., McCanne, S.: Towards an evolvable internet architecture. *SIGCOMM Comput. Commun. Rev.* 35(4), 313–324 (2005)
- [28] Cross-ETP Vision Document, http://www.future-internet.eu/fileadmin/documents/reports/Cross-ETPs_FI_Vision_Document_v1_0.pdf
- [29] Clark, D.D.: The Design Philosophy of the DARPA Internet Protocols, *Proc SIGCOMM 88* (reprinted in *ACM CCR* 25(1), 102-111, 1995). *ACM CCR* 18(4), 106–114 (1988)
- [30] Saltzer, J.H., Reed, D.P., Clark, D.D.: End-To-End Arguments in System Design. *ACM TOCS* 2(4), 277–288 (1984)
- [31] Carpenter, B.: *Architectural Principles of the Internet*, Internet Engineering Task Force (IETF), RFC 1958 (July 1996)
- [32] Krishnamurthy, B.: I know what you will do next summer., *ACM SIGCOMM Computer Communication Review* (Oct. 2010), <http://www2.research.att.com/~bala/papers/ccr10-priv.pdf>
- [33] W3C Workshop on Privacy for Advanced Web APIs 12/13 July 2010, London (2010), <http://www.w3.org/2010/api-privacy-ws/report.html>
- [34] Workshop on Internet Privacy, co-organized by the IAB, W3C, MIT, and ISOC, 8 and 9 December (2010), <http://www.iab.org/about/workshops/privacy/>
- [35] Clark, D., et al.: Towards the Future Internet Architecture, Internet Engineering Task Force (IETF); RFC 1287 (December 1991)
- [36] http://www.iso.org/iso/iso_technical_committee.html?commid=45072
- [37] <http://www.4ward-project.eu/>
- [Alicante] <http://www.ict-alicante.eu/>
- [ANA] <http://www.ana-project.org/>
- [COAST] <http://www.fp7-coast.eu/>
- [COMET] <http://www.comet-project.org/>
- [ECODE] <http://www.ecode-project.eu/>
- [EIFFEL] <http://www.fp7-eiffel.eu/>
- [EULER] <http://www.euler-project.eu/>
- [IoT-A] <http://www.iot-a.eu/>
- [nextMedia] <http://www.fi-nextmedia.eu/>
- [OPTIMIX] <http://www.ict-optimix.eu/>
- [ResumeNet] <http://www.resumenet.eu/>
- [SelfNet] <https://www.ict-selfnet.eu/>
- [TRILOGY] <http://trilogy-project.org/>
- [UniverSelf] <http://www.univerself-project.eu>