

From Internet Architecture Research to Standards

Dimitri Papadimitriou¹, Bernard Sales¹, Piet Demeester², and Theodore Zahariadis³

¹ Alcatel-Lucent Bell Labs, Belgium

{dimitri.papadimitriou,bernard.sales}@alcatel-lucent.com

² Ghent University, Belgium

piet.demeester@intec.ugent.be

³ Synelixis Solutions, Greece

zahariad@synelixis.com

Abstract. Many Internet architectural research initiatives have been undertaken over last twenty years. None of them actually reached their intended goal: the evolution of the Internet architecture is still driven by its protocols not by genuine architectural evolutions. As this approach becomes the main limiting factor of Internet growth and application deployment, this paper proposes an alternative research path starting from the root causes (the progressive depletion of the design principles of the Internet) and motivates the need for a common architectural foundation. For this purpose, it proposes a practical methodology to incubate architectural research results as part of the standardization process.

1 Introduction

The Internet model based on TCP/IP is driven since its inception by a small set of design principles rather than derived from an architecture specification [1]. These principles guided the structure and behavior as well as the relationships between the protocols designed for the Internet. Nowadays, within the Internet community, some argue that changes should be carried on once a major architectural limit is reached (theory of change) and thus the architecture should be designed to enable such changes. Others argue that as long as it works and it is useful for the “majority”, no major changes should be made (theory of utility) as the objective is to keep longevity of the design as much as possible. As a consequence of the theory of utility, the evolution of the Internet is driven by incremental and reactive additions to its protocols or when these protocol extensions are not possible (without changing the fundamental properties of existing Internet Protocols) complement them by means of overlaying protocols. Nevertheless, this approach has already shown its limits. For instance, the design of IP multicast as an IP routing overlay led to limited Internet-wide deployment (even if some have argued that it only enables optimizing capacity consumption without necessarily improving end-user utility). On the other hand, mobile IP (MIP) also designed as an IP network overlay suffers from limited deployment too but it is undoubtedly an essential IP networking functionality to be provided by the Internet.

In this paper, we argue that the debate between the theory of change vs. the theory of utility is reaching its end. Indeed, the representative examples of design decisions provided in Sections 2 aim to explain that the architecture resulting from

independently designed components has already become the limiting factor of Internet growth and deployment of new applications. On the other hand, the incremental and reactive additions to protocols are becoming architecturally complex and thus more and more time consuming; henceforth, this approach has already reached its applicability limit too. This observation leads as explained in Section 3 to rethink holistically the architectural foundations of the Internet itself, and thus, its underlying research process by proposing a “third path” to architecture research. For this purpose, we propose a method can be applied either bottom-up (results drive the model) or top-down (model drives results). The former has also been adopted by the EC Future Internet Reference Architecture (FIArch) Group [2]. This architecture research initiative focuses on key architectural issues and contributes to an EC research roadmap towards a Future Internet Architecture. From the standardization perspective, as the Internet evolution cycle is back to research, the standardization process has also to be reconsidered. As detailed in Section 4, the associated challenges are i) how to best transfer architectural outcomes from research to standard, in particular, by means of the pre-standardization process, ii) how to adapt the standard bodies working methods to accommodate architectural research results, and iii) how to ensure that the architectural results lead to a common baseline.

2 Architectural Model and Analysis

In this section, we provide representative examples of early decisions that drove the design of the current Internet protocols. This non-exhaustive list of examples illustrate perfectly well that the Internet design decisions were taken outside of any holistic architecture albeit critical in the architectural model specification as they impact a large portion of the Internet.

2.1 Architectural Model

Architecture is one of the key elements when engineering complex distributed systems. Surprisingly, it has often been neglected in the context of communication networks design, noticeably to the Internet which remains structured along relatively weak foundations in spite of its ubiquitous deployment [3]. Many definitions of (system) architecture have been formulated over time. In the context of this paper, we refer to the term as “architecture” a set of functions, states, and objects/information together with their behavior, structure, composition, relationships and spatio-temporal distribution. More specifically, the architecture of ICT systems combines three complementary spaces: *functions*: the set of procedures the system performs, their inter-dependencies, and their relationships; *objects/information*: the organization of the data the system processes (input), produces (output) by means of these functions, including their relation and their interactions; and *states*: describing the system behavior as well as the stimuli (condition, events, etc.) that change this behavior. These spaces are modeled using formal techniques including flow block diagrams (functions), object class combined with entity-relation diagrams (objects/information) and finite state machines (behavior). Any “domain” of applicability ranging, e.g., from vending machine to avionic systems, railway signaling system, and large ecosystems will exhibit these three complementary spaces. Hence, they also apply to communication networks in general and to the Internet in particular.

2.2 Architectural Analysis

a) *TCP connection continuity*: from the beginning of the TCP/IP design, it was decided to use the IP address as both network and host identifiers. As TCP has to provide a reliable service, the TCP segments sent by the source to the destination are protected by a checksum that enable to detect whether or not a segment has been corrupted by the network. However, to ensure that a third party host can not inject data traffic over an established TCP connection, the checksum is also performed on the source and destination IP addresses, implying that both addresses shall remain unchanged during the lifetime of the TCP connection. However, the IP address being also a network identifier, the IP address of a mobile host will change together with its attachment point raising issues at the level of the TCP connection continuity. Resolving the latter requires a certain level of decoupling between the identifier of the position of the mobile host in the network graph (network address) from the identifier used for the TCP connection identification purposes.

b) *IP control*: IP forwarding itself is relatively simple but its associated control components are numerous and sometimes overlapping. As a result of the incremental addition of ad-hoc control components, their interactions are becoming more and more complex. As of today, there is simply no systematic design of the IP control functions which in turn causes detrimental effects such as failures, instability, inconsistency between routing and forwarding that can lead to network black holes. Moreover, experience shows that such practice renders the addition of control components exponentially (architecturally) complex leading to overload of existing components. IP routing protocols provides good examples of protocols designed with a reduced set of kernel functions designed with limited flexibility in terms of extension or replacement thus leading to functional overload as soon as expectations on network functionality increases...

c) *Addressing design and routing scaling*: originally, host IP addresses were assigned based on network topological location. Adoption in the mid 90's of dedicated mechanisms to perform address aggregation (called CIDR) was felt sufficient to handle address scaling. Today, conditions to achieve efficient address aggregation and thus relatively small routing tables are not met anymore. This situation is exacerbated by the current Regional Internet Registry (RIR) policy that allocates Provider-Independent (PI) addresses that are not topologically aggregatable; thus, making CIDR ineffective to handle address scaling. The result is that the increase of routing table sizes worsens over time as these prefixes are allocated without taking into account effects on the global routing system. Indeed, routing on PI address prefixes requires additional routing entries in the Internet routing system whereas the "costs" incurred by these additional prefixes, in terms of routing table entries and associated processing overhead, are supported by the global routing system as a whole. Coupled to the increase of the number of routes resulting from site multi-homing (~25% of sites), ISP multi-homing, and inter-domain traffic-engineering, this practice exacerbates the limitations of the Internet routing system. Nowadays, the latter must not only scale with increasing network size and growth but also with a growing set of constraints and functionalities. Hence, routers shall cope with increasing routing table size even if the network itself would not be growing.

d) *Border Gateway Protocol (BGP)*: has been designed to compute and maintain Internet routes between administrative domains. Its route selection algorithm is subject to Path Exploration phenomenon: BGP routers may announce as valid, routes that are affected by a failure and that will be withdrawn shortly later during subsequent routing updates. This phenomenon is (one of) the main reasons for the large number of routing update messages received by inter-domain routers. In turn, path exploration exacerbates inter-domain routing system instability and processing overhead. Both result in delaying the convergence time of BGP routing tables upon topology or policy change. Several mitigation mechanisms exist but practice has shown that the reactive (problem-driven) approach at the origin of the design of these mechanisms does not allow evaluating their potential detrimental effects on the global routing system.

Observations: All these problems could have been avoided or at least mitigated if the Internet was not relying on a minimalistic architectural model. Indeed, a systematic architectural modeling of the system would have i) provided the various possible design options from the beginning and ii) offer to the protocol designer a framework to reason on the role of each of these components and their interactions. Without any architecture model, the components (in particular, the protocols) tend to be designed independently, thus, preventing any holistic approach at design time. Moreover, independent component design does not delimit sufficient condition to achieve global design objectives. For instance, one of the root causes of the Internet scaling resides in the lack of modeling of the global routing system. Indeed, the main choice when designing a routing protocol resides in the selection of the algorithm performing route computation. However, as the routing system is not properly modeled, the impacts of these design choices on the global routing system are almost impossible to evaluate. In contrast, good engineering practices suggest to first model the Internet addressing and the routing system by identifying its architectural components and their relationships. Next, the algorithms for route computation can be designed and their impact on the global routing system can be analyzed and evaluated by using the architectural model. It is to be emphasized here that even if following a systematic and holistic architectural approach does not tell the "right" routing algorithm, this approach can certainly help delimiting what would constitute a suitable algorithm from a functional and behavioral perspective.

What Can We Learn? The Internet architecture is implicitly defined by the concatenation and the superposition of its protocols. In this context, architectural components (in particular, the protocols) tend to be designed independently thus, preventing any holistic approach at design time. Moreover, following the argument of "utility", the evolution of the TCP/IP model is mainly carried out by means of incremental and reactive additions of features to existing protocols relying on a reduced set of kernel functions. This approach has been effectively used since so far but is now progressively reaching objective limits that range from global functional shortcomings and/or performance degradation to maintenance¹ which in turn lead to

¹ Note here that the replacement and/or addition of key architectural components is impossible without changing the properties of the architecture.

serious and costly operational problems. Hence, independent component design does not ensure the sufficient conditions to achieve global design objectives and when achieved lead to detrimental effects. Indeed, reasoning by protocols instead of thinking by functions will ineluctably lead to duplicated functions, conflicting realization of the same function and unforeseen interactions between functions impacting the global system operation. The above examples show that the argument of “utility” is not sufficiently compelling anymore for certain key functions such as mobility, congestion control and routing. On the other hand, the theory of change cannot lead to any significant improvement since there is actually no common architectural baseline (i.e., replacement of an independent component is unlikely to lead to a global architectural change, IPv6 being probably the best example). This corroborates the need for conducting systematic and holistic architectural work in order to provide a proper architectural common baseline for the Internet.

3 Architectural Foundation for the Internet

The disparity of arguments regarding the research path to follow (change vs. utility) is resulting in maintaining the genuine Internet design foundations instead of starting from the root causes: the progressive depletion of the foundational design principles of the Internet. In this section, we argue that the research path to follow is not limited anymore to the selection of the trajectory but the revision of the starting point as determined by these root causes. We contrast the main architectural methods so as to derive a synthetic approach that challenges these foundations.

3.1 Design Principles

Design principles refers to a set of agreed structural and behavioral rules on how an architect/a designer can best structure the various architectural components and describe the fundamental and time invariant laws underlying the working of an engineered artefact. These principles are the corner stone of the Internet design compared to architectures that rely exclusively on modeling. They play a central role in the architecture of the Internet by driving most engineering decision at conception time but also at the operational level. When it comes to the design of the Internet, the formulation of design principles is a fundamental characteristic of the Internet design process that guides the specification of the design model. On the other hand, commonly shared design principles define necessary (but not sufficient) conditions to ensure that objectives are met by the Internet.

Due to their importance, several initiatives have been initiated over last decade that study the evolution of the design principles. Among others, the FIArch initiative has undertaken a systematic analysis of the Internet design principles and their expected evolution [4]. Analytical work on design principles documents the most common design principles of the Internet and put them in perspective of the Internet protocols design and their evolution. These studies aim to identify and to characterize the different design principles that would govern the architecture of the Future Internet.

3.2 Combining Design Principles and Architectural Research

Role of Design Principles: Retrospectively, one of the most advanced architecture for communication networks is known as the OSI Reference Model (OSI RM) standardized in the 80's. Despite its educational value, the protocols architecture derived from the OSI RM did not reach its expectation in terms of deployment. One of the root causes is that the design principles regarding the OSI RM were loosely defined; this practice resulted in lot of protocol misconceptions. Examples include the definition of numerous options in the protocols design that renders interoperability very challenging. This culminates with the creation of two incompatible network layers, one based on connection-oriented and the other based on connectionless (the so-called "CO/CL" debate). While designed at the same time than the OSI RM, the Internet TCP/IP model and its associated protocols are nowadays used ubiquitously as the technologies of the Internet. In contrast to the OSI RM model, the Internet model is driven since its inception by a small set of commonly shared design principles rather than being derived from a formal architecture. For instance, the combination of the "end-to-end" principle with fate sharing suggested the best placement and distribution of functionality taking into account the objective of scaling of IP inter-networks and robustness of TCP at an acceptable cost/performance ratio.

Current Architectural Research: Inspired from [5], current approaches driving architectural research can be subdivided into two categories:

- Driven by the *theory of utility*, this research assumes that the Internet shows longevity and adaptivity thanks to its principles. Its evolution is driven at its "edges" with the expectation to perform capabilities the network alone is unable to provide in particular congestion control (e.g., Explicit Congestion Notification (ECN) and its variants), and traffic-engineering (e.g., multipath-TCP) or by means of overlays (IP multicast, mobile IP but also overlay routing and peer-to-peer fall into this category). It is interesting to observe that independently of the investment and research outcomes, most of these advances have had relatively limited impact on the actual design of the Internet but also its functionality and performance.
- Driven by the *theory of change*, this research assumes that after several iterative cycles of adaptation of architectural components, it becomes more effective to redefine their foundation. Following this approach, the Internet and its design principles are not adapted anymore to address its objectives. The architecture resulting from reactive and incremental improvements to independently designed protocols is already a limiting factor of the Internet growth and the deployment of new applications (at least those that do not directly benefit from capacity addition and/or communication system upgrades). However, in many cases, the result leads to change/replace components as main research objective instead of resolving architectural challenges starting from root cause analysis. A variant of this approach assumes that the Internet can't evolve anymore because under current conditions its design is locked by inflexible systems running processes determined at design time to minimize the cost/performance ratio for a given set of pre-determined functionality. Among prominent efforts falling in this category, we can

mention open-flow, and virtualization but also the more recent software-define/-driven networks (SDN).

Third Path to Architectural Research: Following these observations, we argue that architectural research should follow a "third path" instead of focusing on observable consequences (theory of utility) or its premises (theory of change). This path starts by identifying the actual root causes, i.e., the progressive depletion of the foundational design principles of the Internet and by acknowledging the need for a common architectural foundation relying on a revision of these principles. Indeed, without strong motivations to adapt or to complement the current set of design principles, it is unlikely that the current architectural model of the Internet (TCP/IP model) would undergo significant change(s). If such evidences remain unidentified, the accommodation of new needs either in terms of functionality or performance will simply be realized by the well-known engineering practices residing outside the scope of genuine architectural research work. A representative example is provided by the evolution of the Internet communication stack that leads to reconsider the modularization principle. This principle structures at design time the communication stacks as a linear sequence of modules related by static and invariant bindings. Indeed, when developed CPU and memory were scarce resources and specialization of communication stacks for computers networks lead to a uniform design optimizing the cost/performance ratio at design time. After 30 years of evolution, communication stacks are characterized by: i) the repetition of functionality across multiple layers, such as monitoring modules repeated over multiple layers (which then requires to recombine information in order to be semantically interpretable) and security components each associated to a specific protocol sitting at a given layer (which result into inconsistent response to attacks), which emphasizes the need to define common functional modules; ii) the proliferation of protocol variants (as part of the same layer) all derived from a kernel of common functions/primitives; which emphasizes the need to define generic modules; iii) the limited or even absence of capability for communication stacks to cope with the increasing variability and uncertainty characterizing external events (resulting from the increasing heterogeneity where communication systems proliferate); this observation emphasizes that the functional and even performance objectives to be met by communication systems could vary over time (thus, messages would be processed by variable sequence of functions determined at running time); and iv) the inability to operate under increasingly variable running conditions resulting from the increasing heterogeneity of substrate on top of which communications stacks are performing. Altogether these observations lead to reformulate the modularization principle in order to i) connect functional modules by realization relationships that supply their behavioral specification, ii) distinguish between general and specialized modules (inheritance), and iii) enable dynamic and variable bindings between the various modules such that the sequence of functions performed is determined at running time. In turn, the newly formulated principle provides the means to, e.g., ensure coordinated monitoring operations and account for all security constraints (that comprises robustness, confidentiality and integrity) consistently across all functions performed by the communicating entities.

3.3 Top-down vs. Bottom-up Method

Following Section 3.2, one can postulate that the Future Internet architecture should rely on a revised set of design principles that governs the specification of the generic and specialized components associated to a common architectural model [3]. Starting from various research results (obtained by projects following the "third-path"), the method to specify a common architecture model as defined in Section 2.1, can either be top-down (model drives results) or bottom-up (results drive the model):

- *Top-down method* (see Figure 1): starts (using knowledge from research results) by defining the global architectural level with generic and common specification including function, information and state (step_1). Then these elements are specialized in order to fit the needs of the domain(s) to which they apply (step_2). By specialization we mean here the profiling of certain function and/or information while keeping the generic properties associated to the global level. Finally these specialized elements are translated into system level components (step_3). The challenge here consists in specifying these components from the top so as to produce appropriate building blocks (e.g., protocol components).
- *Bottom-up method* (see Figure 2): starts by exploiting research results and position them as either global (network-level) or local (system-level). In most cases, the corresponding elements are specialized since realized in order to reach architectural objectives that are domain-specific. The challenge with this method then consists in deriving from this set of common and generic components underlying the architecture. Once identified, the result of this step is fed back in order to align the specification of global (network-level) or local (system-level) specific elements. Note there are no actual steps in this method that is characterized by iterative cycles of updates between generic and specialized specification.

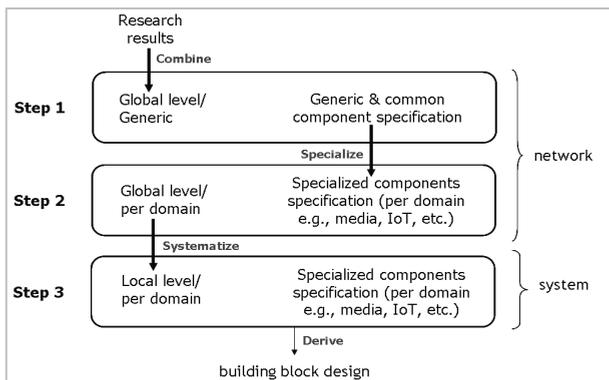


Fig. 1. Top-down method

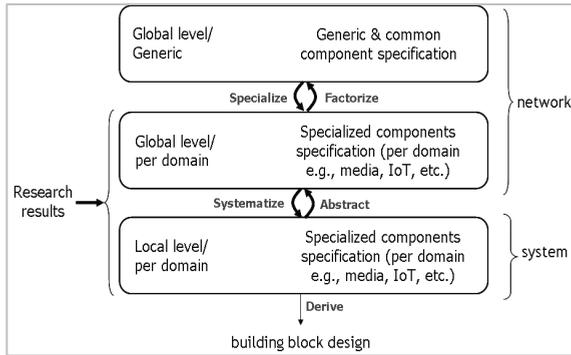


Fig. 2. Bottom-up method

4 Standardization Aspects

As the evolution cycle of the Internet architecture is back to research, the standardization process has also to reconsider its working methods i) to enable best transfer the architectural results obtained from various research efforts to standard, in particular, by means of the pre-standardization process, ii) to accommodate architectural research results, and iii) to ensure that these results altogether lead to a common architectural baseline. Considering that standardization is crucial in the wide adoption of this baseline, this section proposes a methodology to drive its adoption.

4.1 Methodology to Guide Standardization Projects

In some cases, the standardization ecosystem related to a research project is not ready/in place to progress its standardization objectives. In this case, a research-focused standardization phase needs to complement the classical standardization process to feed it with a stream of de-risked ideas that will, if successful, lead to a full standardized solution. For this reason, this phase is generally referred to as the pre-standardization phase. This paper develops a four-step methodology aiming at guiding research projects to identify their standardization needs and to approach them in a systematic way so that the necessary conditions for a successful adoption in standardization are fulfilled:

1. Frame what needs to be standardized (interfaces, etc.) to allow the technology proposed by the project to be interoperable and deployable at large scale. In general, this step implies the identification of an “initial” architecture.
2. Identify the role and impact of standardization bodies on the technology segment targeted by the project. During this step, standardization bodies are categorized as to the role they may (or not) fulfill in the standardization ‘food chain’, i.e., requirements, architecture, solution/protocol, and interoperability and/or testing.

3. Evaluate the need to improve the standardization eco-system to maximize the chance of success. This can be materialized by the creation of a new (pre-) standardization technical committee and, in any case, requires attracting the major stakeholders in the technology segment.
4. Identify the “structuring” dimensions (i.e., what characterizes the standardization objectives trajectory/path) for the proposed technology/system to define a) the criteria to shape the associated standardization target(s) of the projects, b) the necessary conditions to meet for the technology/system to become standardisable. The output of this step is a standardization objectives trajectory to be realized.

In this context, an (initial) architecture enables to systematically enumerate all the interfaces and to formally analyze which of them needs to be standardized for further transfer of the technology/system to marketable products and/or services.

4.2 What to Standardize and Where

The methodology defined in Section 4.1 is generic as it is applicable to any type of ICT research work. When applied to the architectural research work as proposed in this paper:

- Step 1: determine the design principles and the architectural components that should be standardized
- Step 2: architectural work is being conducted by standardization organizations such as 3GPP, BBF, TMF, OIF, DVB and OIPTV. Their role is to drive the packaging of architecture solutions applicable to a given industry segment (e.g., wireline access, 3G mobile system, optical networks, etc.). In this context, these architecture components are reflecting the role of the various systems involved in the solution such as access/edge/core routers, user terminal, and eNodeB. However, the foundational Internet architecture work we propose to conduct is positioned as an upstream activity that will, at the end, feed these existing architecture initiatives. The bodies where the foundational architecture work can be standardized include IRTF/IETF, ITU-T and ETSI. More precisely, the design principles should firstly be proposed to and evaluated together with the Internet Architecture Board (IAB). Indeed, in terms of global reach, the most natural place to model the architecture and its components would certainly be the IRTF. However, IRTF (and IETF) has never considered formal and holistic architecture work as part of its research groups charters. The ITU-T is currently working on related thematic but not yet on components aspects. ETSI is currently hosting several Industrial Specification Groups (ISG), some of them having a Future Internet architecture scope. Moreover, in the context of FP7 and future EC research programs, it would be easier to connect an ETSI ISG to the workforces currently involved in the FIArch Group.
- Step 3: as a result of Step 2, either IRTF needs to be convinced and willing to step into holistic architectural work, or the current ITU-T work program needs to be reinforced/refocused or new work item needs to be launched within ETSI.

- Step 4: it is unlikely that a standardization body will accept to incorporate the architectural aspects of the research output ‘as-is’ in its standardization work program (cf. structuring dimensions of Step 4), mainly because this architectural work needs further validation. As a result, it is proposed to start the work in the pre-standardization mode following either the top-down or the bottom-up method defined in Section 3.3.

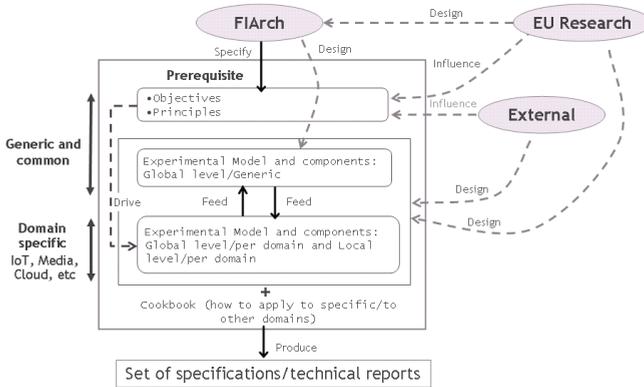


Fig. 3. Application of the bottom-up method

The objectives of the pre-standardization phase will be i) to validate a possible reference architecture model for the future Internet, ii) to give guidelines on alternative models (if applicable), and iii) to explore possible directions for further standardization.

We consider hereafter setting up a potential pre-standardization group (PSG) dedicated to the Future Internet architecture and implementing the iterative bottom-up method as described in Section 3.3. In the context, such a PSG (that could be initiated at, e.g., ETSI or IRTF) will follow the working method proposed in this paper:

- The design objectives and principles provided by the EC FIArch group will be used as input and will drive the specification of the model and components part of the architecture. However, the industry and academia at large will have also a way to influence them in order to ensure their broad acceptance.
- Research projects and academia will contribute to the architecture work in their domain of expertise; these inputs will be used to build an acceptable architecture per key domain, e.g., sensor networks/Internet of Things (IoT), networked-media and data-centers/cloud.
- Common building blocks shared between these domain-specific architectures will be identified; next, the domain-specific architectural components will be aggregated to create the generic building blocks and the relationship between these blocks will be determined.

4.3 Steps toward Architectural (Pre-)Standardization

Standardization work is generally driven by the contributing efforts of its participants. Pre-standardization work on the Future Internet architecture will not depart from this mode of operation. Note that concerning the generic architecture work, the effort will be conducted within the architecture pre-standardization group (PSG). In any case, two distinct scenarios could be envisaged either in the global and/or the domain-specific architecture context:

- First scenario: when a core contribution is submitted to the PSG, it will naturally serve as the fundamental basis for the further discussions. This will happen, for instance, when a team agrees on a given architecture outside the PSG and submits it as input to the pre-standardization process. In this case, contributions from other participants will be in most cases limited to improvements with respect to the core proposal, consensus will be easier to reach and the process will converge quickly to a consolidated architecture reflecting the view of the PSG participants.
- Second scenario: two (or more) competing proposals are brought to the PSG. In this case, reaching consensus within the group becomes more challenging. The way to resolve potential conflicts at the architectural level includes i) the identification of common components (from the competing proposals) fulfilling a completely or partially similar role with respect to the architecture and organize the model accordingly, and ii) the organization of the components that are complementary around a kernel of common components. It is also possible that competing models are actually complementary, leading to an architecture where both models are loosely “interconnected” with a few number of data and/or function relationships depending on the model being specified.

The architecture PSG will be chartered for a limited lifetime. When the global/generic and domain (both local and global) architectural models produced by the PSG are validated and considered as mature enough, the PSG work will have to be transferred to the normal standardization process. The global/generic architecture work can be standardized by the body hosting the PSG. However, concerning the domain-specific (both local and global) architecture work, the standardization body hosting the PSG may not cover all technical domains. As a result, for some domains, the work needs to be reassigned to another standardization body, thus inducing the creation, within the targeted standardization bodies, of one or more architecture working groups/work items to work on a first standardized version of the domain specific architecture. Then, any proposal for new technology, solutions, or protocols will have to be positioned with respect to the existing architecture, including a clear analysis on the impacts of the new proposal on the architecture. If necessary, changes to the existing architecture should be identified and clearly motivated. It should be noted that the architecture will also drive further specification of protocols in order to realize the implementation of its identified interfaces. This new architecture-centric approach is expected to realize in the standardization context the “Third Path to Architectural Research” proposed in this paper (cf. Section 3.2).

5 Conclusion

In this paper, we argue that the debate between architectural research driven by the application of the theory of utility and the theory of change is over. Indeed, neither of these approaches can fundamentally address the limits of the Internet architecture. Instead, we propose that architectural research should follow a "third path" starting by identifying the actual root causes (by adapting the foundational design principles of the Internet such as the modularization principle) and by acknowledging the need for a holistic architecture (instead of (re-)designing protocols independently and expecting that their combination would lead to a consistent architecture at running time). The proposed path will in turn also partially impact how the necessary standardization work is to be organized and conducted. Indeed, the design principles and the modeling part of the architecture need to be standardized to ensure its adoption at the international level. Following this path, the chartering of the new work item to define, e.g., new protocol, will need to be not only "problem-driven" but also "architecture driven". It is also anticipated that, resulting from the current wave of Future Internet research projects, the pre-standardization work will become more and more relevant with a mix of architecture- and technology-driven work items. As such, this is an opportunity since this nascent pre-standardization ecosystem can be seen as a laboratory to learn how to introduce an "architecture-driven" dimension in the Internet standardization working method.

Open Access. This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

1. Carpenter, B.: Architectural Principles of the Internet. IETF, RFC 1958 (June 1996)
2. European Commission, Future Internet Reference Architecture Group (FIArch), http://ec.europa.eu/information_society/activities/foi/research/fiarch/
3. Papadimitriou, D., Sales, B.: A path towards strong architectural foundation for the internet design. In: 2nd Future Network Technologies Workshop, ETSI, Sophia Antipolis, France, September 26-27 (2011)
4. FIArch Group Report v1.0. Future Internet Architecture-Design Principles (January 2012) (in press)
5. Clark, D.: Designs that last. In: FP7 EIFFEL Think Tank, Athens, Greece, October 6-7 (2009)
6. Zahariadis, T., et al.: Towards a Future Internet Architecture. In: Domingue, J., et al. (eds.) FIA 2011. LNCS, vol. 6656, pp. 7–18. Springer, Heidelberg (2011)