PURSUIT
Publish Subscribe Internet Technology
FP7-INFSO-ICT-257217

TR 12-0001

From Information-centric Networking to Information-centric Architectures

Document Properties:

Title of Contract	Publish Subscribe Internet Technology
Acronym	PURSUIT
Contract Number	FP7-INFSO-ICT 257217
Start date of the project	1.9.2010
Duration	30 months, until 28.2.2013
Document Title:	From Information-centric Networking to Information-centric Architectures
Date of preparation	2012-04-03
Author(s)	Dirk Trossen (UCAM)
Responsible of the deliverable	Dirk Trossen Email: dirk.trossen@cl.cam.ac.uk
Target Dissemination Level:	PU
Status of the Document:	Final
Version	1.0
Document location	http://www.fp7-pursuit.eu/
Project web site	http://www.fp7-pursuit.eu/

Table of Contents

Executive Summary

Information-centric networking has been attracting increasing attention in the research community with a variety of architectural proposals and technology solutions. While such progress is laudable, there has been little debate as to what the real potential of this area is in the larger space of networks and systems. A skeptical view at the promise of improving the large-scale dissemination of information might point at the continuous efforts of industry and academia alike to improve speed and reach of the current Internet. So what is the real potential for information-centric networking? This report provides a viewpoint to this question by creating a link to the way we design systems at large, going beyond mere operational improvements in parts of today's Internet.

1. Introduction

The Internet has been exceeding all expectations with respect to its capability to cope with the modern demands for information dissemination. The developments in serving multimedia content over the Internet provide staggering examples: 3 billion videos are watched on YouTube every day with 48 hours uploaded every minute [1]; on average, more than 250 million photos are uploaded on Facebook per day [2]; in the UK, the BBC's iPlayer platform is used on average 68 minutes per week by each viewer with 3.1 million daily viewers in September 2010 [3]; and in all this, mobile usage just started to take off with mobile video traffic expected to exceed 50 percent of total mobile data for the first time in 2011 [4].

Furthermore, other areas of digital information are accelerating the trend of ever-increasing information dissemination in the Internet. Several national-level 'open government initiatives' [5][6] make a huge amount of data available from government departments for the usage by citizens and organizations. And forecasts, e.g., in [4][7], for areas such as sensor networks, supply chain management, health, retail and many others expect the supply of information into our Internet to grow significantly within the next years.

Alongside this development, the area of *information*- or *content-centric networking* (called ICN in the following) has been increasingly attracting attention in the wider research community, fueled by research efforts in various parts of the world. Originated from initial efforts in 2006 [8], many other research efforts, such as [9][10][11], have taken up the promise of this area with the aim to improve the way networks will be able to move around information at large scale. While progress has been made at the level of specific design choices, hard evidence on the expected improvements is still inconclusive. Furthermore, since the inception of the area some five years ago, the current Internet has been evolving tremendously in order to cope with the ever-increasing supply of digital information. These observations lead us to believe that

> *The initial claim of ICN to improve the performance of content delivery can no more be a convincing argument for such radical paradigm change nor does it fully reflect its true potential!*

Instead, we see the true potential for ICN in leading us to a new way of designing layered information-centric communication architectures. Insights derived from ICN will allow for a flexible assembly of strictly defined layers, each of which can be optimized across the full space that is provided by the computational and communication resources that are available. We believe that this promise can be combined with recent findings, which define layering as a process of optimization decomposition [12], into a rigorous framework for system design that is truly novel as well as powerful.

It is this system level impact that raises the discussion on information-centric networking above the particular performance improvements it might bring for parts of its deployment and it places ICN at the heart of a discussion on how to design distributed systems at large. This is a discussion we intend to start in this article, while leaving much of the needed evidence as future work.

For this, we structure our presentation as follows. In Section 2, we first outline aspects on layered architectures that arise from the understanding that has been created within the community through the past 30 plus years. We utilize the IP-based Internet as a case study that points us at potential lessons to be learned with respect to layered architecture in general. We then highlight in Section 3 some of the changed constraints that exist today as compared to the time of inception of the current Internet, these changes being fueled by technological advances. It is Section 4 in which we outline the potential insight that information-centric networking might bring in the wider problem field of layered architectures. We then extend this particular connection to system design beyond communication systems in Section 5.

2. On Layered Architectures

Let us begin with outlining the major insights into large-scale system design, insights that follow from 30 plus years of communication system design that has led to the development of today's Internet[1]. We refer to the situation in the current Internet as one that can inform us about lessons to be learned when thinking about designing an evolved, yet far larger scaled successor but also when thinking beyond communication systems.

One crucial aspect that has been recognized in many research efforts [13][12][14] is the *robust-yet-fragile* nature of the Internet as a complex system. It is the robustness against perturbations, for which protocols in individual layers have been designed, while it is the fragility with respect to unforeseen perturbations that can lead to, often catastrophic, failures [14]. The layered nature of the Internet is seen at the heart of this robust-yet-fragile nature that we can observe. Hence, a rigorous understanding of layered architectures is crucial for any progress in designing complex networked systems, pioneered in systems thinking by Day [15], in control theoretic terms by Doyle et al. [12], and in the socio-economic dimension by Clark et al. [16]. It is here that we see a connection to the proposition of information-centric networking. Before delving deeper into this connection, however, let us elaborate on the rigorous understanding of layered architectures.

In a distributed system, an extreme heterogeneity exists in the space of content and services, being distributed over a likewise heterogeneous set of resources at the infrastructural level. Layered architectures and the set of protocols defined within each layer attempt to homogenize this view on content and services as well as resources through strictly defined interfaces, effectively *constraining through a well-defined protocol*. These exposed interfaces provide a clearly defined feedback through which the individual protocols can be optimized against a set of well-defined constraints, providing the robustness that can be observed against perturbations for which the protocols are designed. It is this set of constraints that provide a convenient basis for formalizing the functionality to be realized within the constraints, with the possibility to eventually lead to a mathematical foundation for its solution. Hence, individual layers provide a *strictness* within which individual optimization problems can be expressed and eventually be solved[2]. Through the set of exposed interfaces, a process of information exposure (and hiding) is defined, manifesting a well-defined information flow across layers. It is this information flow across layers, which enables the creation of economic markets through defined *information asymmetries* [18].

Individual layers provide the robustness against the constraints for which the individual protocols are designed, while the *process of layering* itself modularizes the individual protocols along well-defined functional spaces. In other words, *flexibility* is introduced in the presence of optimized resource utilization in each individual layer, with feedback mechanisms across layers ensuring an overall robustness of the architecture. These feedback mechanisms utilize the information flow across the defined interface of each protocol. Such modularization (which is present in a layered architecture) suggests a plug-and-play nature of the overall architecture[3]. It requires true flexibility in the ability to modularize in order to fully come to fruition, i.e., to achieve optimization decomposition across a large space of resource utilization, as envisioned in [12]. We will return to this issue of flexibility when discussing the situation in today's Internet.

---

[1] It is worthwhile pointing out that much of the system-level developments in the Internet have been post-rationalizations of the engineering intuitions that drove the design decisions. It could be argued that it is this combination of top-down rationalization and bottom-up intuition to which the success of the Internet can be attributed.

[2] TCP FAST [17] is an example for forward-engineering a solution through optimization theory.

[3] It is this modularization that [12] describes as *deconstraining the constraints* introduced by each optimization process that is represented by each (modular) protocol. Clark et al. [16] outline this deconstraining aspect as addressing the tussles between the modularized spaces of functionality; these tussles often existing in the socio-economic rather than purely technological space.

Viewing layered architectures as an approach to deconstrain the constraints of individual layers through flexible layering directly connects system design to the optimization over a space of resources that are at the disposal of a system designer. Within the context of designing large-scale communication systems, the resources available are those of *computation* over and *storage* of information as well as the *transfer* of information through communication links. Utilizing the full space of optimization across these dimensions of resources is at the heart of any solution design. This utilization of the full optimization potential equates to the widest possible enablement of markets through flexibly enabling information asymmetries. The constraints that limit this optimization potential are *technological* and *economic* ones; we return in Section 3 to the changes in these constraints since the inception of the IP-centric design of the current Internet. Another issue is the lack of an integrated theory that combines the service-oriented application world with that of the optimization-oriented view of computational and communication resources in the infrastructure. As a result, design largely decomposes the larger problem a priori, leading to inefficiencies and fragility in the long run due to the emergence of perturbations that diverge from the expectation of the designer. We believe that information-centric networking has the potential to provide new insight into this issue, which we will briefly address in Section 5.

Layering as a principle is well recognized and utilized in the current Internet, seemingly aligning the current Internet with the framework of optimization decomposition. However, we point out two issues that limit the current Internet's ability for optimization.

The first one relates to the underlying constraints that determine the utilization of resources in the Internet. The Internet's original design was perceived at a time when computational resources were scarce; computing centers were limited in availability and storage was very expensive in terms of price per bit. These economic and technological constraints favored a design, which assumed information to exist behind dedicated *portals*, each of which being identified with a routable address. This focus on *location* of information only allows for a limited exploitation of computational resources once they become more widely available. However, as we show in Section 4, these constraints have vastly shifted towards an almost ubiquitous availability of such resources, all of which are rather inefficiently utilized in a design that requires an early binding to location (expressed in a process of address space allocation and management). Improvement attempts in today's Internet, such as content delivery networks, as well as efforts in the space of grid and cloud computing promise to remove this limitation. These attempts, however, are only moderately successful since the underlying substrate for reaching the various computing clusters still relies on a location-oriented dissemination of information. As a conclusion, we assert that *any approach to bind location at an early stage of communication limits the ability to fully utilize the space for optimization*.

The second limitation in today's Internet is the layering model itself. Although it is generally the foundation for deconstraining the individual protocol constraints, the rigidity of the IP layer model, i.e., the introduction of a particular assembly of layers, limits the ability to flexibly modularize a wide variety of constraints that are possibly introduced by the individual layers after the original design. Such wider variety of constraints occurs, for instance, when moving into environments that were not envisioned in the original design. Such change in constraints leads to a possible fragility. We can observe examples for this fragility all over the Internet, such as through the introduction of network address translation (NAT) or firewall technology as an additional set of constraints to the originally designed full end-to-end connectivity between IP end hosts. Another example is that of virtual private networks (VPNs), introducing additional security constraints into the communication platform. Moreover, the world of wireless communication has brought about an entire range of new constraints for which the given assembly of layers, in particular the routing and transport layer, proves to be too rigid. While the emergence of these new constraints might pose difficulties for the existing architecture, such as breaking end-to-end connectivity through NATs, it also points to possible solutions for future architectures to come, e.g., by introducing local addressing only. In today's world, however, solutions to these additional constraints are introduced as shim or middlebox solutions, leading to the observed fragility in the overall behavior of the network. These examples are all motivated by a changing context of optimization constraints, which are then bolted onto an inflexible architectural foundation, namely a rigid layering system[4]. The very existence of these solutions shows the desire for such design flexibility towards changing optimization contexts.

This leads us to assert that *such desire for flexibility needs to be properly captured in an approach to system design rather than needing to rely on kludges in an increasingly inflexible design*. We recognize that this assertion directly connects to the discussion on *ossification of the Internet* [34] through the increasing complexity caused by various patches and extensions to original design decisions. While our assertion partially follows the often stated need for a *clean slate*, it more importantly asserts that the system design approach taken for such new start is more important than the new start itself.

Day [15] accommodates this desire with the introduction of a *recursive layer model*. In the following, we build on this thought while extending it through a connection to information-centric networking. But let us first expand on the changing constraints that cause so many headaches in today's systems.

3. On Changed Constraints

As discussed in Section 2, fragility in large-scale systems arises from perturbations that have not been foreseen by the original system designer. While these perturbations mainly refer to unforeseen elements that arise during the operation of the system, such as an uncharacteristic packet loss due to the introduction of a novel transmission system, we focus in this section on a longer-term economic and technological trend that has fundamentally changed the set of constraints in which today's Internet operates. It is this trend that leads us to a new starting point for large-scale communication systems; a starting point that is deeply rooted in the ideas of information-centric networking.

The first trend is that of the *availability of storage and computing resources*. At the time of conceiving the Internet's fundamental design, storage and computing resources were highly scarce, with the notion of personal computing still to be invented. The original Arpanet connected dedicated computing facilities for the purpose of resource sharing; a resource that could hardly be replicated in many places due to the enormous price point that they represented. Hence, transmitting bits between these dedicated computing centers represented an alternative in favor of replicating the resource itself at a different place; remote access was key. This relation between *bandwidth* (for remote access) and *storage & computing* (for the performance of computational tasks) has significantly changed during the past 30 or more years. Fueled by Moore's Law, computing and storage has become significantly cheaper, in particular compared to the alternative of transferring the information through communication systems. For illustration purposes, Figure 1 depicts the development of hard disk prices per gigabyte over 15 years, with similar trends to hold for computing resources in terms of *flops per dollar*. But not only has storage and computing become cheap, it more so has become almost ubiquitously available, in particular in the developed world[5] [19].

It is this technological trend that is at the heart of propositions in today's Internet, such as web proxies and content delivery networks, to allow for a more flexible optimization between the resources of *communication* and *storage*[6]. But this trend is also at the core of information-centric networking, recognizing that the WHAT of a communication is likely to exist in many more places than the WHO that originally disseminated the information.

---

[4] The fact that we refer to these solutions as 'cross-layer violations' shows the unpreparedness of our current design approach to accommodate such new sets of constraints!

[5] Even the computing resources in simple feature phones in the developing world represent a significant increase of computing power compared to the early 70s!

[6] We need to recognize that it requires kludges like request interception and DNS redirection for integrating such desires for optimization into the daily operation of a system that was not originally designed for it.

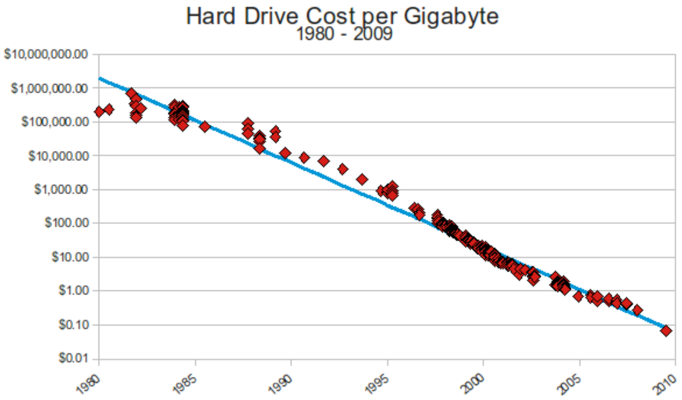

Figure 1: Costs per Gigabyte Storage [20]

The ubiquitous availability of resources has also changed the attitude of many users that utilize the Internet. While real-time communication is still a major part of the daily usage of the Internet, services that utilizes distributed resources for computation and storage are on a constant rise. Bittorrent-like large-scale dissemination of files, being stored only in parts within the resources made available by the users of the system, has found widespread adoption in the Internet. In addition, large-scale distributed computing platforms for complex event processing are at the heart of any modern enterprise infrastructure with contemporary counterparts in the academic world such as the SETI project [21] that utilizes volunteers' contributions in the search for extraterrestrial intelligence. This has led to a changing set of abstractions that is utilized on the top of an infrastructure that is heavily location-based (by virtue of location endpoint identifiers that only change, if at all, at mid- to long-term time scales). This new set of abstractions that is utilized by these new usages of the Internet is focused on information dissemination rather than connecting endpoints[7].

Another area that significantly changes the constraints for possibly realizing communication systems is that of *wireless communication*. While in its infancy at the time of designing today's Internet, it has nowadays become a dominant form of accessing communication resources worldwide; in fact, for many users in the developing and developed world alike, wireless access has become the only means of accessing the Internet (often even for the first time). This has driven frequency allocation of the wireless spectrum to be changed very much in favor of wireless communication, with recent developments of freeing up parts of the spectrum for the ever-increasing demand for wireless Internet access [23]. Wireless communication, however, is in its nature significantly different from wireline communication in that a notion of a-priori location of an endpoint hardly exists[8]. This leads to a plethora of solutions for mobile communication that are insufficiently integrated and supported in an all-IP network. Solutions such as [24] address this situation by introduction cognitive techniques in the assembly of networks, while approaches such as Haggle [25] propose to overcome this deficiency through a gateway approach to networking, bridging various architectural approaches through Haggle nodes. An encompassing architectural framework that properly integrates wireless resources as yet another optimization dimension into the communication system, however, is still missing.

The take-away of our discussion here is that apart from perturbations to the system that arise from unforeseen constraints in the operations of the system, there is a range of longer-term trends that have significantly changed the constraints for designing a communication system. It is this longer-term trend that we see better addressed by information-centric networking.

---

[7] The authors in [22] argue that this level of abstractions at the HTTP level constitutes a possible new waist of the Internet with a proven, albeit increasingly complex foundation of protocols.

[8] Treating wireless technology similar to its fixed counterpart, the Ethernet, reconciles that difference, albeit at the cost of inefficiencies in utilizing the wireless nature of the medium (such as for mobile scenarios, increased resilience, etc.).

4. From ICN Towards Information-Centric Architectures

Let us now turn our attention to information-centric networking. While the changing trends, outlined in Section 3, have been referred to in other works, e.g., [26], it is the connection to the design of layered architectures together with our desire to fully utilize the potential for optimization that we see at the heart of the game-changing aspect of ICN.

An inspiration to think about layered architectures in relation to information-centric networking comes from efforts on *meta-reasoning architectures*, and more specifically in the architecture defined by Chappell and Sloman [27]. In this work, the authors capture the meta-reasoning and deliberative processes in human cognition. They devise a functional architecture that integrates the *reactive processes* that we humans so quickly and subconsciously execute in everyday activities with the more conscious processes of *goal-oriented reasoning*. Most operations of this system are executed as 'fast path' operations, while slower meta-reasoning and deliberation only occurs when an attention filter threshold is exceeded, i.e., when attention by the more complex reasoning processes is deemed necessary.

Let us now apply the Chappell/Sloman model to networked systems by focusing on *solving problems* throughout such systems. Paraphrasing [28], problems involve "*a collection of information that*" an implementation *"can use to decide what to do"*, which is to implement a problem solution. The basic information required includes states, goals, and actions. Problems have a nature of spawning sub-problems or relate to other problems. Hence, they can be *concatenating* or *inclusive*. As an illustration, consider as a first example that of link-level fragmentation of information into individual chunks that can be suitably transferred over a medium such as Ethernet. For this, further consider a node A that wishes to send a video of length *N* via a forwarding graph to node B. Assume at least one of the links along the graph being restricted in its maximum length of information that can be transferred, denoted by *M*. If *N*>*M*, fragmentation needs to occur on all of the links with such restriction. Hence, the problem of sending the video from A to B needs to be extended to solve the individual sub-problems of sending individual fragments of size *M* on restricted links. Figure 2 visualizes this inclusive nature of problems in this case.

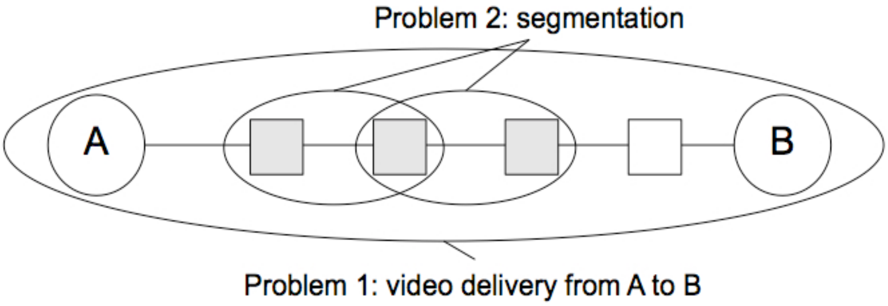

Figure 2: Fragmentation Problem Example

We can extend this problem to a larger network. Apart from fragmentation, other sub-problems are likely in need to be solved. When considering video delivery along a multicast tree, the individual forwarding decisions at the branching points are examples for other sub-problems. Caching the video (in the case of non-real-time distribution) could be another problem to be solved. Adjustments of the operations could occur through changing operational parameters, such as the maximum transfer unit used across several fragmentation solutions, adjusting spanning tree parameters or branching point positions, and adjusting possible caching strategies.

When solving such problems through networked systems, *computation* and *storage* resources are utilized in elements such as end hosts, routers, switches, caches, and even links between these elements. The communication system facilitates such resource usage through the exchange of any relevant information. This places *information* at the core, which constitutes a significant departure from the IP paradigm, in which the location of endpoints is the main concept (with opaque bits being transmitted between these endpoints). Such focus on information and its dissemination for solving individual problems defines a *problem-centric approach* to building systems with deliberative and reactive processes as shown in Figure 3.

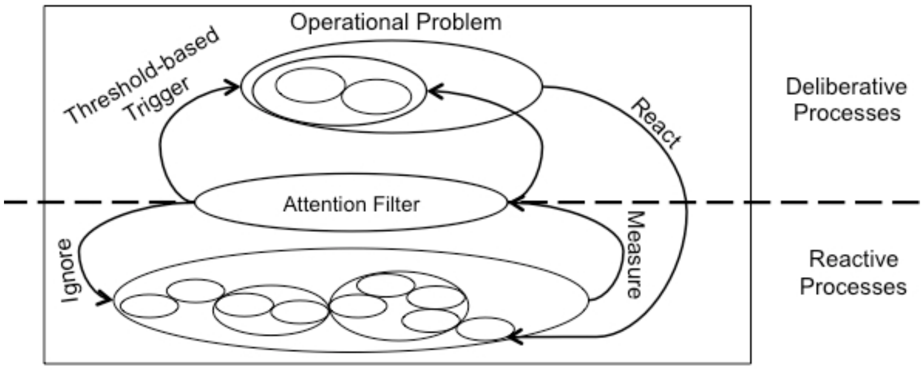

Figure 3: Problem Solving through Meta-Reasoning

At the heart of such problem-centric approach is the process of *mediation* between interest and availability of information, in order to fully utilize computational and communication resources. This concept of mediation connects the information-centric concepts in programming languages and operating systems with the constrained, location-dependent view of resources at the infrastructural level. In other words, the process of mediation ensures a temporal and spatial decoupling of information dissemination, which in turn enables a greater potential for fully utilizing the resources at the disposal of the system. From an optimization viewpoint, we assert that *a systems approach that operates on graphs of information*[9] *with a late binding to location at which the computation over this graph is going to happen, enables the full potential for optimization over the entire space of communication and computational resources.*

This assertion highlights the computation over a graph of information, where the role of the communication system is that of optimally disseminating information through a late binding to a location for an entity that can optimally contribute to the solution. The late binding provides us with a key ingredient to optimize against the constraints in individual layers of the architecture. The desired flexibility across layers is captured through modularizing information-centric functionality within a distributed system. Given our problem-centric view to design, a recursive nature of such flexible layering comes as an almost natural requirement.

This translation of our problem-centric view onto communication systems that revolve around information dissemination leads us to a set of desired system properties that we see at the heart of a layered system architecture that enables the full potential for optimization. Firstly, a system ought to provide means to *manipulate structured information flows* for computational purposes. Secondly, a system ought to clearly *modularize individual problems*. And thirdly, a system should *modularize across computational problems* in order to assemble these towards larger ones. From a layered architecture point of view, this leads to a recursive nature of layering around well-defined computational problems (each of which can be optimized within its defined constraints). With this, we can recognize a resemblance with the work of Day [15], albeit applied in an information-centric context. In other words, our properties change the structure of Day's recursive layers from interpreting any communication as an inter-process communication (IPC) between possibly remote machines to a structure in which computational problems directly operate on a graph of information, as defined by our property 1. For that, the communication system exposes a service model that enables the manipulation of information structures within a computational context. It is the communication system that ensures the

---

[9] It is important to understand that our chosen information graph abstraction is only one possible choice. The memory concept in [29] presents an abstraction to application programmers that can be seen as an alternative to our graph concept.

separation of the WHAT from the WHO of an information exchange by virtue of a late binding of location to the particular information within a given structure.

Such late binding to location is ensured by a set of core functions in each layer, implementing the process of mediation with late-binding delivery in each layer. It is this modularity that enables the optimization of these functions within a given architectural decomposition of the larger problem space. While each of the layers provides a strict framework for solving an optimization problem, flexibility is provided through our third property that envisions a (recursive) modularization across individual computational problems towards a larger one. Hence, this flexibility deconstrains the constraints introduced by individual (strictly defined) protocols within each problem solution. More importantly, this plug-and-play nature enables a problem-specific modularization and is therefore a departure from the current Internet model that introduces a common and fixed layering structure throughout all problems.

Following these desired system properties, a next step is that of deriving tenets for designing large-scale distributed systems; an exercise we will leave for a broader contribution on system architecture. However, we see the area of information-centric networking as one that can lead us to this set of design tenets for layered architectures that can fundamentally change the way we design scalable distributed networked systems. Information-centrism and flexibility in assembling computational problem solutions that utilize information dissemination are at the heart of this new design approach.

A wider discussion beyond this editorial is required to qualitatively argue for the properties and tenets that, as we assert, enable an increased optimization potential compared to the system design approach that is at the core of today's Internet. However, we have first indications of such increased potential in concrete areas of optimization, such as [30][31]. Furthermore, a first artifact is available that realizes the desired system properties in a novel node architecture design, demonstrated in a high-speed data network [32]. While this cannot be seen as conclusive evidence, it is a promising start for a wider architectural debate.

5. Towards the Big Picture

Throughout this article, it has been our ambition to embed information-centric networking into a context that is more than content dissemination and an evolutionary improvement of the current Internet. Instead, we asserted a connection between ICN and system design in general. This embeds ICN into a design approach that aims at exploiting a wider range of optimization potential than what is given to us with the current Internet. It is this placement of ICN in the context of system design, however, where we see the true contribution of ICN by informing the wider system design community on what could be the principles of layered system design. In other words, we strongly believe that ICN provides us with a starting point to formulate a more general approach to large-scale system design; an approach that provides a strong connection between engineering intuition and experience and the growing set of mathematical models and tools for robust control and optimization. We assert that, as an outcome of such work, we will see an increased understanding on how the design of large-scale distributed systems can be fundamentally rigorous, can accommodate a wide range of heterogeneous environments, can react to unforeseen perturbation with reconfiguration, and can, in all this, optimally utilize the resources at its dispense.

But the biggest ambition of all is the potential insight that such work could have on the design of systems in general, not limited to communication artifacts such as the Internet. As so compellingly argued by Crowcroft in [33], we see tremendous potential from the lessons learned on rigorous system design in other areas of engineering and science. This cross-fertilization of insights into design in general already exists in selected examples, such as biology. What is missing, however, is a unifying *theory of design* that cuts across these disciplines. We firmly believe that our work on ICN provides first insights into such theory.

6. Acknowledgements

The design considerations presented in this paper have been the result of intensive discussions among the researchers in the EU FP7 PURSUIT project (http://www.fp7-pursuit.eu). We specifically thank John Doyle, John Wroclawski, Janne Rihhijarvi, Karen Sollins, Jon Crowcroft, and George Polyzos for their feedback as well as Dana Pavel for her insight into meta-reasoning architectures and the connection to ICN.

7. References

[1] YouTube Statistics, at http://www.youtube.com/t/press_statistics, 2011

[2] Facebook Statistics, at http://www.facebook.com/press/info.php?statistics, 2011

[3] BBC iPlayer, at http://www.bbc.com/iplayer, 2011

[4] Cisco Inc, “Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2010–2015”, February 2011

[5] “Opening Up Government”, at http://data.gov.uk/, 2011

[6] “Open Government Initiative”, at http://www.whitehouse.gov/open, 2011

[7] IBM, “Managing the Explosion of Information”, Report, 2008

[8] Van Jacobson, “A new way to look at networking”, at http://video.google.com/videoplay?docid=-6972678839686672840, 2006

[9] T. Koponen et al., “A Data-Oriented Network Architecture”, ACM SIGCOMM, 2007

[10] V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs, R. Braynard, “Networking Named Context“, Proceedings of ACM CoNext conference, December 2009

[11] D. Trossen (ed.), “Conceptual Architecture: Principles, Patterns and Sub-Component Descriptions”, available at http://www.fp7-pursuit.eu, 2011

[12] Chiang M, Low SH, Calderbank AR, Doyle JC (2007) Layering as optimization de-composition: A mathematical theory of architecture. Proc IEEE 95:52–56.

[13] J. Doyle, D. L. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, W. Willinger, “The robust yet fragile nature of the Internet”, Proceedings of the National Academy of Science, 2005

[14] W. Willinger, J. Doyle (2004), Robustness and the Internet: Design and Evolution. In *Robust design: A Repertoire of Biological, Ecological, and Engineering Case Studies*, E. Jen (Editor), Oxford University Press

[15] Day, J., “Patterns in Network Architecture - A Return to Fundamentals”, Prentice Hall, 2008.

[16] D. Clark, J. Wroclawski, K. R. Sollins, R. Braden, “Tussle in Cyberspace: Defining Tomorrow’s Internet,” in IEEE/ACM Transactions on Networking, Vol. 13, No. 3, 2005

[17] Wei, David X.; Jin, Cheng; Low, Steven H. and Hegde, Sanjay (2006). "FAST TCP: motivation, architecture, algorithms, performance". IEEE/ACM Trans. on Networking 14 (6): 1246–1259.

[18] B. C. Greenwald, J. E. Stiglitz, “Externalities in Economies with Imperfect Information and Incomplete Markets”, Quarterly Journal of Economics, no. 90, May 1986

[19] IDC, “The Diverse and Exploding Digital Universe”, EMC-sponsored whitepaper, 2008

[20] Cost per gigabyte, at http://www.mkomo.com/cost-per-gigabyte, 2011

[21] SETI@home, at http://setiathome.berkeley.edu/, 2011

[22] L. Popa, A. Ghodsi, I. Stoica, “HTTP as the Narrow Waist of the Future Internet”, Proc. of HotNets workshop, 2010

[23] M. Mahoney, “Spectrum of Issues, MIT Tech. Review, at http://www.technologyreview.com/communications/26568/, 2010

[24] R.W. Thomas, L.A. DaSilva, A.B. MacKenzie, "Cognitive networks", in: Proceedings of the First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks, Baltimore, MD, USA, November 8–11, 2005

[25] J. Scott, J. Crowcroft, P. Hui, C. Diot, “Haggle: a networking architecture designed around mobile users”, In Proceedings of IFIP WONS 2006

[26] D. Trossen, M. Sarela, K. Sollins, “Arguments for an Information-Centric Internetworking Architecture”, ACM Computer Communication Review, April 2010

[27] J. Chappell, A. Sloman, "Natural and artificial meta-configured altricial information-processing systems", Invited contribution to a special issue of The International Journal of Unconventional Computing, Vol 2, Issue 3, 2007

[28] S. J. Russell, P. Norvig, "Artificial Intelligence: A Modern Approach", 2nd Edition, Pearson Educ., 1998

[29] I. Piumarta, "An association-based model of dynamic behaviour", Proceedings of Workshop on Free Composition at ECOOP, 2011

[30] C. Chau, R. Gibbens, T. G. Griffin, "Towards a Unified Theory of Policy-Based Routing", Proceedings of INFOCOM, 2006

[31] P. Jokela, A. Zahemszky, S. Arianfar, P. Nikander, C. Esteve, "LIPSIN: Line speed Publish/Subscribe Inter-Networking", ACM SIGCOMM, August 2009

[32] Blackadder prototype, at https://github.com/georgeparisis/blackadder, 2010

[33] J. Crowcroft, "FIE: Future Internet Enervation", ACM SIGCOMM Computer Communication Review, Vol. 40 Issue 3, July 2010

[34] Stanford Notes, "Ossification of the Internet", at http://www.scs.stanford.edu/nyu/04sp/notes/l23.pdf