Will ICN Fly?

A Case Study in Technology Adoption

Börje Ohlman* and Kostas Pentikousis^

* Ericsson Research, S-164 80 Stockholm, Sweden

Email: Borje.Ohlman@ericsson.com

^ EICT, EUREF-Campus Haus 13, Torgauer Strasse 12-15, 10829 Berlin, Germany

Email: k.pentikousis@eict.de

Abstract

Information-centric Networking (ICN) proposes a paradigm shift in the way we create global networks. ICN introduces new network primitives which operate on named data objects, regardless on which host they reside, therefore representing a clear departure from the host-centric paradigm of the current Internet. The peer-reviewed ICN literature documents several system design advantages over host-centric networking. While ICN research is currently addressing several technical aspects that can facilitate future deployment for a range of scenarios, we also need to consider other aspects in technology adoption beyond technical superiority. In this paper we examine previous technological developments and pinpoint the economic, operational, societal, and policy-related aspects that should be paid attention to as we move forward in ICN deployment from isolated testbeds into real networks over the next decade.

Introduction

How does novel research get from the design board and isolated testbed deployments into widespread real-network use? This paper reviews the adoption of recent (and historical) technical innovations and aims to identify the characteristics of widely deployed innovations in contrast to those that failed to get traction with real-world users. This exercise is of course valuable for a range of modern networking technologies and protocols standardized at the IETF and elsewhere. But, we are also interested in how could Information-centric Networking (ICN) real-world adoption benefit from the lessons given by previous deployment experiences.

Our goal with this paper is to contribute to the general discussion on what “makes or breaks” a new technology. For this discussion we have chosen ICN as an example technology. The first, obvious, reason is that ICN is our current field of research and therefore such an exploration can assist our own work. The second reason is that ICN has currently reached a level of maturity that makes deployment possible in the mid-term. One example of a practical ICN implementation was demonstrated by Cisco at SIGCOMM 2013: an NDN router with a line-rate throughput of 20 Gbit/s. We hope that this paper can help us, the ICN research community, learn from previous equally ambitious efforts to shake up the hitherto technological status quo and shape up entire future industries.

Although we use ICN as our example technology in this paper we do not assume that the reader is an expert in ICN. Interested readers can see [1, 2] and the references therein for an introduction to ICN. On the contrary, this position paper aims at stimulating the discussion about technology adoption, in general, while considering the transition from a host-centric paradigm to ICN. In particular, we hope that this paper does open up a wider discussion in the community and the IAB about how can ICN succeed in the real world.

The remainder of this paper is organized as follows. The following section reviews technology adoption in networking, operating systems, and telecommunications. We then examine ICN deployment opportunities based on current real-world problems. We conclude this paper with a short discussion on the two areas where ICN may have a significant edge over IP, namely media distribution and the Internet of Things (IoT).

Technology Adoption in Retrospect

This section reviews trends in the adoption of technology in the computing and telecommunications spheres.

"Good enough", at the Right Time

It is often argued that "good enough", not perfect, technology should be in the minds of developers [3, 4]. Indeed, a good-enough technology, that addresses a pertinent problem space and is introduced at an opportune moment, can come to dominate developments for years. Said technology may lack solid answers to a range of issues at the time of its introduction. Yet, experience indicates that it is often preferable to take a solution which has the important aspects right, but is not fully specified, closer to its users and develop it from there, rather than wait until all technical aspects are spelled out in detail.

Well-known landmark examples of technologies that were "good enough" for a wide adoption and thus won over technically superior alternatives include VHS vs. Betamax , and DOS/Windows vs. Mac in the 80s and 90s. VHS is also an example of "Content is King". What users primarily wanted was content, i.e. rental movies. As VHS was first out [5, 6], most content was available for VHS. This meant that most users invested in VHS equipment, which in turn lead content producers to prioritize content releases on VHS initiating a virtuous cycle for VHS and a death spiral for Betamax.

Another example is the deployment of Linux. Although the GNU project had been working for some time on several applications, common libraries, and the Hurd kernel, the project had yet to achieve a critical mass of user adoption. A critical piece, i.e. an OS kernel for inexpensive x386 systems, was missing [7]. The availability of the Linux kernel, designed particularly for this category of PCs, enabled the creation of a critical user mass and served as the foundation for a multi-billion dollar industry over the years. It is worth recalling that at the time of its release, several authoritative figures in the OS field were highly critical of the Linux kernel, pinpointing to several faults in a very vocal manner [7]. Availability, however, as well as the formation of a supportive community, made sustainable development and deployment possible.

Understanding the problem space, seeing the solution space

Another important aspect understanding the process of technology adaptation is getting a comprehensive understanding of the problem space and the solution space, and then being able to see the intersection. An example of this is the development of modern digital mobile telephony. One of the earliest developments of mobile telephony was the Nordic Mobile Telephony (NMT) system developed in the Nordic countries. The Nordic countries are quite large in area but with populations that are sparsely and unevenly distributed throughout their geography. This provided a strong motivation to develop mobile telephony. At the same time, limited population sizes provided a small potential user base, which led system designers to pool resource in order to bring down the development costs. Hence, NMT was developed as an international project and cross-border roaming became one of the fundamental design requirements early on [24].

In contrast, in the U.S., mobile telephony was introduced by non-cooperating regional operators. This created fragmented network coverage with limited roaming capabilities. Since national roaming was perceived as a significant problem, international roaming seemed simply out of reach. We believe that for many in the U.S. telecommunications industry the only feasible solution to make global mobile telephony roaming possible seemed to be to have one operator spanning the globe. Arguably, worldwide connectivity and real-time communications were key drivers for satellite telephony systems [8]. In the end, however, terrestrial cellular telecommunications based on global roaming provided a less expensive [9], more practical alternative with better incentives for all stakeholders.

Another example in this category is the development of the iPhone, which eventually changed the direction of the entire mobile device industry. Steve Jobs had a clear vision of the device he wanted to build. He was convinced that people would want it even though the focus (or "early adopter") group that was calling for it was not well-defined. In the designer's mind, plainly put, it was all pretty straightforward: People like the Internet— and everyone would like it in their pocket. This answer lies fundamentally at the intersection of understanding the problem and seeing the solution. Of course, the iPhone would not become reality without steering a large amount of resources in the same focused direction. This new direction, indeed for the entire industry, would not have been successful without capitalizing on the market forces behind the operator needs for a competitive advantage (cf. AT&T vs. Verizon in terms of network coverage and quality the US circa 2007) [10].

Market Dominance and Oligopolies

The traditional telecommunications industry has a history starting with traditional national monopolies for telephony and oligopolies in the mobile market where licensed spectrum is an insurmountable entry barrier [11]. In the mobile telephony sphere, telecommunication systems tend to be very large and complex [9], and are developed at a fast pace when compared to earlier and fixed telecommunication technologies. This was possible due to the commitment of a few market-leading manufacturers together with incumbent operators that in unison advocated the technologies to be adopted [12]. Notably, this can be seen in the evolution of 3GPP 3G/4G/LTE that was pushed through standards bodies [13] not so much because of disproportionate technical supremacy [14] over competing solutions, but due to the market dominance of the players who advocated it [15].

The Role of Total Cost of Ownership and Expanding Markets

The failure of Quality of Service (QoS)-related technologies (e.g. IntServ, DiffServ and RSVP) in the Internet can serve as an example where we believe that Total Cost of Ownership (TCO) may have posed a major stumbling block. While, in theory, these technologies could provide for significant savings on resource usage, in practice, there were no business cases to support it. To get those savings on resources the provider had to sell a significantly more complex service to the user and manage a significantly more complex network. For most users, this meant a less attractive service, which thus would have had to be offered at a lower price than what a flat rate provider would offer to be accepted. The cost for a flat rate provider using overprovisioning turns out to be very low at a constantly growing market. For instance, assume that the provider needs to double the capacity in the network every 12 months in order to meet new user demand. The cost of running the network at 50% utilization can be approximated with the extra capital cost of doing the required network infrastructure investments 12 months ahead of time. Of course, this equation provides different answers in a steady-state market.

Summary

We identified four factors that play an important role in getting wide market adoption. Namely, we argued that "good enough" solutions can often win over more advanced technologies which require further work before they could be deployed. We noted that code availability and community building can play a crucial role in succeeding in the marketplace. Second, we observe that understanding the intersection between the problem and solution space can provide a great advantage in the medium term and shape entire industries in the long term (such as the cellular infrastructure network and mobile devices industries).Thirdly, we point out that forming strong alliances that share common interests can foster technology adoption much faster. Finally, we note that TCO and market growth prospects are at least as important as increased resource efficiency that a particular technology can offer. It goes without saying that combinations of these factors also exist, making the corresponding technology even more formidable.

We will now move on to discuss how the future of future technologies can be determined by these factors. We will then go on to use the example of ICN to see how this can be applied.

Windows of Opportunities for ICN

This section discusses briefly scenarios that the current host-centric network architecture has problems coping with. In effect, these examples spell out opportunities for ICN deployment.

Video content explosion

There is no sign that the ongoing explosion of video content should taper off. On the contrary, network operators, content creators and providers are genuinely concerned about how to keep up with the traffic growth [16]. This provides an opportunity for technologies that enable caching at the network edges, including inside home networks and end devices, which can serve content locally and directly.

Prosumers

Digital content prosumerism is well established today [17]. The trend towards increasingly more content being created and served from the edges of the network challenges the archetypical topologies of the host-centric Internet which is configured to support point-to-point communication and distribution of content from media servers. Distributing live video (e.g. from a natural disaster site) using a single Android device to a worldwide audience is, at the very least, a true challenge in today's host-centric network architecture. This opens a window of opportunity for technologies that can handle this type of flash crowd events better that the host-centric paradigm.

Denial of Service

As end devices become more capable and access links increase in capacity, the possibility to mount large-scale denial of service (DoS) attacks increases. New technologies that do not allow pushing traffic towards specific end-points in the network but require active content requests by the receivers are advantageous as they can robustly mitigate this type of attacks.

Privacy

To have your personal information stored in the servers of global corporations has its privacy issues [18-21], as the recent NSA scandal has emphatically brought to the fore. Technologies, which allow users to share information in social networks and/or other communities without having to store the information in a centralized infrastructure, have the potential of giving users better control over who can access their information.

ICN through the Crystal Ball

Much of the current focus of ICN research is directed towards two areas, namely, media distribution (primarily video) and the so-called Internet of Things (IoT) [22]. The case for media distribution is quite straightforward. The inherent caching and multicast capabilities of ICN make it a natural successor to today's CDN solutions *if* there will be a future network where this functionality will be integrated in the basic network service. For IoT, a key feature of ICN when compared to today's IP networking is that ICN does not require end-to-end (e2e) connectivity. The inherent ICN caching infrastructure lends itself well to the required store-and-forward functionality enabling battery powered devices to remain offline most of the time and ensures that needless operations, including delivering the same piece of data twice, are avoided.

Having said that, and pointed at two specific issues that need to be resolved in a future (ICN-based) Internet, one should remember that the key feature of ICN is that it offers the opportunity to resolve many pain points at once while rolling out a new clean(er) architecture. For any advantage we can claim with ICN (caching, multicast, multihoming, better mobility support, flash crowd robustness, etc.) it is easy to claim that a specific fix to the current Internet can solve one of these issues. One proposal for dealing with a number of these issues for HTTP traffic is idICN [23]. But doing individual fixes for each and every one of these issues might create an architectural mess and a network that will be very difficult to manage.

Conclusion

This paper provided a retrospect on the factors that affect technology adoption and discussed several examples of successful technology adoption. We then identified opportunities defined by the current network infrastructure paradigm for a new technology like ICN in the near future. In particular, we looked at the possibilities for ICN technology to get widespread adoption and identified the factors beyond technical excellence that could play a significant role towards the transition to an ICN-based Internet. As part of our ongoing work we are looking into migration paths and deployment cases for ICN that may serve as the foundation for establishing a critical mass for this technology in the real world.

References

1. Ahlgren, Bengt, et al. "A survey of information-centric networking." Communications Magazine, IEEE 50.7 (2012): 26-36.
2. Pentikousis, Kostas, et al. "Information-centric networking." Communications Magazine, IEEE 50.7 (2012): 22-25.
3. Yourdon, Edward. "When good enough software is best." *Software, IEEE* 12.3 (1995): 79-81.
4. Sandhu, Ravi. "Good-enough security." *Internet Computing, IEEE* 7.1 (2003): 66-68.
5. Park, Sangin. "Quantitative analysis of network externalities in competing technologies: The VCR case." *Review of Economics and Statistics* 86.4 (2004): 937-945.
6. Carlsson, Bo, and Rikard Stankiewicz. "On the nature, function and composition of technological systems." *Journal of evolutionary economics* 1.2 (1991): 93-118.
7. DiBona, Chris, and Sam Ockman. *Open sources: Voices from the open source revolution*. O'Reilly, 2008.
8. Pratt, Stephen R., et al. "An operational and performance overview of the IRIDIUM low earth orbit satellite system." *Communications Surveys & Tutorials, IEEE* 2.2 (1999): 2-10.
9. Bi, Qi, G. L. Zysman, and Hank Menkes. "Wireless mobile communications at the start of the 21st century." *Communications Magazine, IEEE* 39.1 (2001): 110-116.
10. Zhu, Ting, Hongju Liu, and Pradeep Chintagunta. "Wireless Carriers' Exclusive Handset Arrangements: An Empirical Look at the iPhone." *Available at SSRN 1962799* (2011).
11. Paterson, Nancy. "Walled gardens: the new shape of the public Internet." *Proceedings of the 2012 iConference*. ACM, 2012.
12. van de Kaa, Geerten. "Flourish or Perish: The Influence of Interorganizational Networks and Flexibility on Technology Success in Winner-take-all Markets." *Proceedings of DRUID 2012*.

13. Mcqueen, Darren. "The momentum behind LTE adoption [sGPP LTE]." *Communications Magazine, IEEE* 47.2 (2009): 44-45.
14. David, Martín-Sacristán, et al. "On the way towards fourth-generation mobile: 3GPP LTE and LTE-advanced." *EURASIP Journal on Wireless Communications and Networking* 2009 (2009).
15. Murphree, Michael, J. P. Shim, and Kevin Love. "National approaches to technology standards for mobile telecommunications." *Wireless Telecommunications Symposium (WTS), 2013*. IEEE, 2013.
16. Weber, Bertrand. "Content Delivery Networks: Market Overview and Technology Innovations." *Media Networks: Architectures, Applications, and Standards* (2012): 209.
17. Bruns, Axel. "From prosumer to produser: Understanding user-led content creation." In *Transforming Audiences 2009*, 3-4 Sep, 2009, London.
18. Liu, Yabing, et al. "Analyzing Facebook privacy settings: User expectations vs. reality." *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. ACM, 2011.
19. Lipford, Heather Richter, Andrew Besmer, and Jason Watson. "Understanding Privacy Settings in Facebook with an Audience View." *UPSEC* 8 (2008): 1-8.
20. Tene, Omer. "What Google knows: Privacy and internet search engines." (2008).
21. Cranor, Lorrie Faith, and Norman Sadeh. "A Shortage of Privacy Engineers." *Security & Privacy, IEEE* 11.2 (2013): 77-79.
22. Dannewitz, Christian, et al. "Scenarios and research issues for a network of information." *4th International Mobile Multimedia Communications Conference*. 2008.
23. Fayazbakhsh, S., et al., "Less pain, most of the gain: incrementally deployable ICN". *Proceedings of ACM SIGCOMM 2013*. ACM, p. 147-158.
24. Meurling, J., Jeans, R., *The Ugly Duckling*, Ericsson Mobile Communications AB, 1997, ISBN 91 630 5452 3