Bridge Networking Research and Internet Standardization: Case Study on Mobile Traffic Offloading and IPv6 Transition Technologies

Aaron Yi Ding[†*], Jouni Korhonen[‡], Teemu Savolainen[§], Markku Kojo[*], Sasu Tarkoma[*†], Jon Crowcroft[†]

[†]University of Cambridge [*]University of Helsinki [‡]Renesas Mobile [§]Nokia

Abstract. The gap between networking research communities and Internet standardization organizations (SDOs) has been growing over the years, which has drawn attention from both academic and industrial sides due to its detrimental impact. The reason behind this widening gap is complex and typically beyond the mere technology ground. In this position paper we share our perspectives toward this challenge based on our hands-on experience obtained from joint projects with universities and companies. We highlight the lessons learned, covering both successful and under-performed cases, and further suggest viable solutions to bridge the gap between networking research and Internet standardization, aiming to promote and maximize the outcome of such collaborative endeavours.

1 Introduction

The proliferation of Internet has boosted the development and deployment of Internet protocols worldwide. The growth generates strong demands for innovative protocol design that further promotes the networking research and standardization. However, a growing gap between research communities and standardization organizations (SDOs) has cast shadows to the sustainable development of the Internet. The problem at large involves multiple stakeholders and depends on various factors on top of the technical aspects, thus making it hard to analyse and tackle with dedicated solutions.

To guide the Internet protocol design, RFC 5218 [1] summarizes a set of key factors and requirements for successful global deployment. Meanwhile, except for general discussions on the transfer of R&D to commercial products in IT business, there are few dedicated studies or investigations on how to bridge the existing gap between networking research and Internet standardization. In seeking feasible solutions to bridge such widening gap, we present our case study on two trendy topics: mobile traffic offloading and IPv6 transition technologies. Through the tight collaboration between academic researchers and industrial experts participating in our joint projects, we have achieved good results in terms of academic publications and standardization contributions. Along the progress, we also learned many lessons that, we believe, are valuable for the Internet community.

In this paper, we make the following contributions:

1. We identify the main challenges and opportunities for the collaboration between networking research and Internet standardization, i.e., Internet Engineering Task Force (IETF), and break down the problem domain based on our hands-on experience obtained from joint projects.
2. We share observations obtained from our research-industry collaborative projects [2–5]. Our case study focuses on two specific topics, mobile traffic offloading and IPv6 transition technologies, which recently have drawn great attention and investment from both research communities and SDOs. Our project [2] on the two topic areas has yielded in total 5 published IETF RFCs, 2 documents in RFC editor's queue that will soon become RFCs, 6 Internet Drafts that may never proceed to RFC status, and a number of academic publications.
3. We further offer our insights on the lessons learned from the successful and under-performed cases, and propose viable suggestions.

The rest of the paper is organized as follows. In Section 2 we analyse the challenges and opportunities to aggregate efforts from research communities and Internet SDOs. We present our case studies in Section 3 and Section 4, covering mobile traffic offloading and IPv6 transition technologies, respectively. We share our lessons and suggestions in Section 5 and conclude in Section 6.

2 Challenges and Opportunities

To enhance the technology transfer in IT industry, there are earlier works investigating how to promote the collaboration for commercial success [6–8]. In recent years, we are facing unique challenges for the widening gap between networking research communities and Internet SDOs that require a dedicated analysis. Based on our hands-on experience through joint projects, we identify several challenges including technical and non-technical aspects.

From the technical viewpoint, we observe three visible challenges:

- ◇ Tools and Expertise: The tools and platforms used by researchers and industrial experts often vary from each other. As dedicated training is required for professionals to enter another new field, such specific tool set forms a crucial barrier in terms of skill and expertise.
- ◇ Process and Target: The scientific publishing and standards development follow different processes. The first visible difference is the review process adopted by research community and SDOs such as IETF. Such variation directly leads to the timing concern. For instance, the time required to publish a scientific paper ranges from three months to few years depending on the venue. Whereas for IETF standardization, it can take several years for a proposal to be approved as an RFC. Another challenging issue is the target, i.e.,

where and how to deliver contributions. In particular, academic researchers are well aware of where and how to publish scientific papers but quite unaware of how to contribute to SDOs, and such knowledge, in general, is not open nor publicly shared.
- ◇ Project Organization: Organizing and running projects vary across academia and industry. To guarantee successful deployment, industrial projects often target at running code and validation in live environments. Whereas academic projects typically focus more on the theoretic formation and solid verification through modelling, simulation, and/or prototyping. Such difference between the agile industrial style and the progressive development of research forms another challenge.

Concerning the non-technical aspect, we observe several challenges that are more subtle and harder to quantify. We provide here a non-exclusive list according to our own experience.

- ◇ Career Development: A crucial hindrance for academic researchers to contribute to SDOs is the potential negative impact on their career development and promotion. Although too harsh to confess, researchers in general hardly reach senior positions such as professorships by conducting standardization work. At the same time, industrial experts face similar challenges for promotion when they mainly publish scientific articles.
- ◇ Performance Evaluation: There is clear gap between the performance metrics of evaluation used in academia and industry. For academic researchers, the major metrics include scientific publication, citation, thesis supervision, teaching, and managing research projects. In general, exploring novel ideas that lead to technology breakthrough is highly valued in the academic track. Meanwhile, the key metrics for industrial professionals include patent, standardization, product development, system and network deployment, value-added service design, and managing industrial R&D projects. Comparing to the academic track, incremental and deployable contributions that produce commercial profits are the first concern for standardization professionals.
- ◇ Financial Support: Lack of financial support and understanding from direct superiors or institute leaders is another inevitable challenge that directly affects professionals' motivation to attend meetings or conferences to learn from and collaborate with other experts.
- ◇ Fear: There is a subtle overestimation for the difficulty of the work outside the own field, i.e., the unrealistic mental fear that the expected failure rate will be higher than the ground truth.
- ◇ Disrespect: Underneath the fear, we observe also the disrespect in a very deceptive manner. It is not unusual to encounter statement at either conference seminars or standardization meetings that the work from research community is unrealistic and incomprehensible, while industrial work being purely money-driven and short sighted.

Although facing several challenges, we believe there are benefits and opportunities behind a close collaboration between the research community and SDOs.

For researchers, there are two major incentives: 1) extending research work to standardization increases the impact and visibility of academic contribution; 2) exposing research to real world problems helps extend the horizon of researchers and results in better outcome that benefit the overall Internet society. For standardization experts, there are also two appealing opportunities: 1) integrating the research helps improve the robustness of industrial R&D work, which further guarantees successful deployment in real market; 2) the power of research also helps overcome the limitation of testing in relatively small-scale industrial labs.

Concerning joint benefits, researchers need testing facilities and real world problems that can be supported by industrial partners. For example, companies can provide experimental environment such as access to core network facilities, usage traces, and challenging problems that are regarded as rare and valuable input for academic research. Meanwhile, industrial companies focusing on standardization need strong analytical tools and expertise from research community to help overcome scale limitations in their testbed settings and to verify their solutions before pushing products to the competitive massive markets. For both sides, combing the analytical studies with large scale experimental investigation of real world problems can yield outcome with great significance. Take our joint work for instance, we managed to produce publications and standardization documents that promote the Internet development. Such joint efforts also create a collaborative environment for industrial experts and academic researchers to work together and learn from each other.

3 Case Study: Mobile Traffic Offloading

Due to the fast increase of mobile data traffic volume generated by bandwidth-hungry smartphone applications, cellular operators are forced to explore various possibilities to offload data traffic away from their core networks. As a part of our project lasting from 2010 to 2013, we investigated this domain from both research and standardization point of view.

For research, we achieved good results [9–13]. The area was topical and mobile operators were searching for a solution, which gave our research a justification to be carried out in the first place. The approach we took was slightly controversial within telecom circles. The decision was intentional, since we thought the approach taken by 3GPP was an overkill and had too much functionality on other layers than IP.

For standardization, we proposed our protocol design for IPv4 traffic offloading [14] to IETF intended for IETF standardization primarily in the Multiple Interface (MIF) Working Group [15]. We also briefly touched the base in 3GPP SA2 with our ideas. For an average IETF and likewise a 3GPP proposal we were rather well prepared. We had running code for two solution approaches, one based on DHCPv6 and the other based on IPv6 Neighbor Discovery protocol. This included implementations in Linux kernel, commercial smartphones and a live network system prototype.

At that time, being able to demonstrate selective IPv6-based network controlled offloading between WLAN and a live cellular network, even while roaming, was not possible for that many IETF participants or even 3GPP delegates.

We summarize the outcome in Table 1.

Table 1: Project Outcome for Mobile Traffic Offloading

	Research	Standardization
Contributions	5 publications in good venues	1 IETF Internet draft, 1 3GPP SA2 Change Request
Performance	Very good	Under-performed

If we look into RFC 5218 metrics on what makes a successful protocol, we could argue that our solution met most of the criteria. However, whether the offloading solutions had *Positive Net Value (Meet a Real Need)* (RFC 5218 Section 2.1.1.) is not straightforward. The solutions definitely might have issues with the existing business models of both operators and vendors. Furthermore, one can argue that the offloading solutions have challenges to claim to have *Good Technical Design* (RFC 5218 Section 2.1.7.) since they, for example, mix IPv4 into IPv6 only protocols and introduce routing style of functionality into IPv6 Neighbor Discovery protocol that is typically hard to justify in IETF. When it comes to 3GPP side of the coin, then our offloading solution, intentionally, tried to make a step away from some long standing design principles how to differentiate IP flows.

In general, this track yields promising scientific results, but standardization effort can be regarded as an under-performed case. Frankly, the authors were aware that the offloading topic is challenging to drive in standardization, specifically in IETF working groups that are not already familiar with cellular network technologies (such as the mobility working groups). Almost the same could be said about 3GPP. Proposing a solution that does not fit to an existing mindset is usually a failure to being with, unless there is a really strong hype or desperate need for such a solution. These were not met, specifically when having a *research stamp* on the solution proposal.

4 Case Study: IPv6 Transition Technologies

During the transitional phase toward IPv6, it is crucial to guarantee the IPv4-IPv6 interoperability for the smooth IPv6 adoption. Among various proposals, the once criticized Network Address Translation (NAT) is gaining a positive role in such transition to bridge the gap between two incompatible IP versions. In our project, we investigated the domain of IPv6 transition technologies by focusing on the mechanisms for the NAT discovery and learning of the IPv6 prefix used for protocol translation in an access network.

For research, we conducted the first extensive comparison study of all solutions in the domain and shared our first-hand implementation experience in real networks [16, 17]. Our study reveals the potential pitfalls that should be considered and offers an empirical basis for evaluating competing mechanisms in the transitional phase.

For standardization, we proposed two competing protocol designs [18, 19] to IETF for standardization primarily in the Behavior Engineering for Hindrance Avoidance (BEHAVE) Working Group [20]. We have running code for both solutions including GNU C library implementation and EDNS0 patches for DNS BIND9. Our experimental analysis has been accepted as Informational RFC. Based on the community feedback and our feasibility test, we recommended the heuristic discovery approach that has been adopted in the IETF transition toolbox for the IPv6 Internet and is pending publication as Standards Track RFC.

We highlight the outcome in Table 2.

Table 2: Project Outcome for IPv6 Transition Technologies

	Research	Standardization
Contributions	1 publication in good venue	2 RFCs and 1 Internet Draft
Performance	Good	Success

Referring to RFC 5218, it is clear that our proposals meet all the initial success factors. When we first proposed our solutions to IETF, there were more than 10 candidate solutions [16]. For the design strength, both of our proposals entered the final round to compete for the final standard mechanism. As shown in Table 3, the EDNS0 approach provides more functionality and efficiency but suffers from the limitation especially concerning its impact on other entities in the network. Although the heuristic discovery offers only moderate adaptation, the main reason for it to excel as the final standard comes from its extensibility (particularly the security extension support) and minimum impact on the infrastructure.

Table 3: Comparison of IPv6 Prefix Discovery and Learning Solutions

Solution	Active detection	Adaptation to changes	Transparent to network	Host system changes	Secure learning
Heuristic discovery	Yes	Moderate	No if DNSSEC	No	With DNSSEC
EDNS0 option	Yes	Fast	No	Yes	No

In general, this track achieved the goal in terms of scientific publication and standardization, hence being a successful case.

5 Lessons and Suggestions

Through our collaborative projects investigating mobile traffic offloading and IPv6 transition technologies, we have learned several lessons. We list the key findings for the under-performed case of mobile traffic offloading:

- ◇ The offloading problem is well comprehended in 3GPP but not in the IETF community which results in extra work to educate and convince IETF professionals why it is an important issue for standardization, in order to clear the resistance. This prolongs the process and complicates our joint work in which researchers and standardization experts are subjected to various restrictions such as time commitment and contract duration. On the other hand 3GPP had already selected their preferred *telecom style* of approaching the offloading topic and trying to convince them to revert to something that basically needs no new work or equipment was not welcomed. Furthermore, in our case, the authors' priority and interest gradually shift away leaving limited resources to complete the work and sell the solution.
- ◇ Research results offer good incentive to bring the work forward to standardization. However, careful considerations are needed to distinguish which part of work is meaningful and suitable as proposed standards. In IETF, the Multiple Interface (MIF) working group [15] we chose is very challenging to push our work and lead the discussions since we started immediately with a solution. It might have been wiser to start first with requirements, and then possibly proceed to MIF. This could have built support in the community at first and possibly made a case needing a solution.
- ◇ From the technical perspective, it is unclear how much room is left for adding extra IPv4 support into the future IPv6 Internet. Our proposal is therefore controversial due to the integration of IPv4 functionality into the IPv6 protocol, i.e., IPv6 neighbor discovery. Although being an incremental solution, our protocol design also affects multiple entities in the existing infrastructure including the networking stacks on mobile hosts and access routers. This results in the resistance from the IETF community and restrains other vendors from implementing our proposal.

For the successful case on IPv6 transition technologies, we share our reflections that can be helpful for the future development.

- ◇ A clearly defined problem leads to acceptance from the standardization community. It also enables ISPs and vendors to decompose the problem and further deploy the proposal addressing the problem.
- ◇ Solutions that can be easily implemented and are deployable will gain essential support from the community and stand out from competing solutions. In our case, we propose two solutions that both enter the final round. The solution that introduces less impact on existing operation and infrastructure wins the final seat.
- ◇ Solid research and well-analysed results increase the acceptance rate and provide the community more input to judge the value of the work and the relevance to become a standard.

To highlight the linkage between standardization and research, we summarize our observation in Table 4. One interesting and surprising finding is on the openness aspect where SDOs have always been making documents publicly accessible, but for networking research, the open-access procedure starts only recently. Meanwhile, the research community has been actively contributing to standardization in terms of implementation, e.g., there is often at least one implementation from academic institutes for standard proposals, normally done by PhD researchers.

Table 4: Linkage between SDO and Academia concerning success and impact

	Standardization	Research
Evaluation Metrics for Success	Approved RFCs Implementations Large scale deployment	Scientific publications Citation Founding new areas
External Impact (Society and Economic)	Open-source/licensed code release Open access reference Service integration/enhancement	Public policy/law amendment Society awareness improvement Tech-transfer via start-up company

On top of the lessons from our case study, we further generalize a set of suggestions based on our experience across research and standardization over the past 20+ years.

1. Focus and Partnership in Joint Projects: One good example is the Trilogy project [21] with its Multipath TCP proposal as a huge success in terms of both academic publishing and standardization. Research institutes such as University College London contributed not only a significant part of thinking, simulation, and prototyping but took very active role also in standardization. With the support from EU partners, Nokia and Apple, the proposal is gradually integrated into commercial OS versions. A key to such success is the patience and investment in time, which can span very long periods (5 to 10 years). Meanwhile, comparing to the large projects, smaller projects offer the advantage that its risk management is easier and the steering of project progress is more agile. With dedicated participants, small projects can yield good results especially in terms of individual contribution level. There are good examples following the principles we summarize, including the WiBrA [2], IoT [3] and HAT projects [22, 23].
2. Personnel Transfer: This targets at junior level professionals such as PhD students. There are numerous successful stories where graduates take on the ideas from their PhD research and realize them through standardized solutions and/or commercial platforms in the companies they work for. This process can be fast but depends on the right person and as well as the size of problem (smaller than #1 in general). There are also intermediate ways of doing this such as sending students to industry through internship, although the idea transfer are even smaller. For industry, the PhD hires at top Internet

companies represent good investment. At the same time, the recent trend for industrial experts moving back to academia also improves the situation in terms of sharing expertise and experience.

3. Understanding the Difference: If a problem is well defined, being real in existing network but relatively easy to fix, the standardization goes smoother. However, if a problem is hard and the fixing process is complex, the authors would have to prepare to write much better problem descriptions, do much more footwork in convincing people of the problem and feasibility of a solution, and generally more ready to spend time in selling the solution. For instance, to get such solution through IETF, it is important to invest time to talk with professionals and do sufficient background work. Simply presenting the problem and solution rarely leads anywhere. This is different from typical research perspective where problem itself is more valuable than the complexity of deployment and solutions. Another difference is the time commitment. While research papers can take 1+ year to appear in conference and/or journal, in an SDO, making an idea to publish may take much longer. In our IPv6 heuristic discovery draft, it took from September 2010 to tentative September 2013, with 20 revisions leading to the upcoming RFC, i.e., 21 document revisions and three years in total. Such time frame can be too long for a researcher to be involved, regarding the fact that 1-2 year project funding is common nowadays, which may demotivate the researchers even to start the process.
4. Organizational Support: For industry experts, one typical barrier to attend academic conferences is the gap between the money investment and payback in the form of profitable product. The same applies to researchers who finds little value to contribute standardization. There are fortunately organizations like Nokia Research Center and Deutsche Telekom Innovation Labs that stand between academic and industry organizations, and can do both. Such organizational support also increases the awareness across domains. It is often that industrial experts are not aware of the latest research results and even so, do not know where to obtain such information, similar to the situation where researchers being unaware of the latest standardization progress and the existing solutions on the product level. The organizational support can form a channel providing pointers, and help solve the challenge of career promotion for both researchers and industrial experts.

6 Conclusion

This position paper presents our case study exploring how to bridge networking research and Internet standardization. We believe there is great potential behind comprehensive collaboration of research and SDO communities. We wish our lessons and suggestions can shed light on this challenging domain and further promote the sustainable development of the Internet.

References

1. D. Thaler and B. Aboba: What Makes for a Successful Protocol? Internet RFCs, ISSN 2070-1721, RFC 5218, 2008.
2. Wireless Broadband Access (WiBrA) Project: `http://www.cs.helsinki.fi/group/wibra/`
3. TiViT Internet of Things (IoT) Project: `http://www.iot.fi`
4. S. Tarkoma, and H. Ailisto: The Internet of Things Program: The Finnish Perspective. *IEEE Communications Magazine*, 51(3):10–11, March 2013.
5. European Institute of Innovation and Technology (EIT) ICT Labs: `http://www.eitictlabs.eu/home/`
6. A. Reisman: Technology Transfer: A Taxonomic View. *Journal of Technology Transfer*, Summer-Fall:31–36, 1989.
7. G. A. Prosser: The Role of Incentives in the Deployment of Technologies from Cooperative R&D. *Journal of Technology Transfer*, 20(2):13–17, 1995.
8. J. Lane: Understanding Technology Transfer. *Assistive Technology*, 11(1):5–19, 1999.
9. A. Y. Ding, B. Han, Y. Xiao, P. Hui, A. Srinivasan, M. Kojo, and S. Tarkoma: Enabling Energy-Aware Collaborative Mobile Data Offloading for Smartphones. In *Proceedings of IEEE SECON 2013*.
10. J. Korhonen, T. Savolainen, A. Y. Ding, and M. Kojo: Toward Network Controlled IP Traffic Offloading. *IEEE Communications Magazine*, 51(3):96–102, March 2013.
11. A. Y. Ding, P. Hui, M. Kojo, and S. Tarkoma: Enable Energy-Aware Mobile Data Offloading for Smartphones through Vertical Collaboration. In *Proceedings of ACM SIGCOMM CoNEXT PhD Student Workshop 2012*.
12. A. Y. Ding, J. Korhonen, P. Hui, T. Savolainen, S. Tarkoma, and M. Kojo: NAO: A Framework to Enable Efficient Mobile Offloading. In *Proceedings of ACM Middleware PDT Workshop 2011*.
13. J. Korhonen, T. Savolainen, G. Wolfner, and J. Laganier: Evolving 3GPP Bearer Model Towards Multiple IPv6 Prefixes and Next-Hop Routers. *Springer Telecommunication Systems Journal - Special Issue on Mobility Management for Flat Networks*. (Accepted, to appear)
14. J. Korhonen, T. Savolainen, and A. Y. Ding: Controlling Traffic Offloading Using Neighbor Discovery Protocol. IETF Internet Draft, 2012, work-in-progress.
15. IETF Multiple Interfaces (MIF) Working Group: `http://tools.ietf.org/wg/mif/`
16. A. Y. Ding, T. Savolainen, J. Korhonen, and M. Kojo: Speeding up IPv6 Transition: Discovering NAT64 and Learning Prefix for IPv6 Address Synthesis. In *Proceedings of IEEE ICC 2012*.
17. J. Korhonen, and T. Savolainen: Analysis of Solution Proposals for Hosts to Learn NAT64 Prefix. publication pending as IETF Informational RFC.
18. T. Savolainen, J. Korhonen, and D. Wing: Discovery of IPv6 Prefix Used for IPv6 Address Synthesis. publication pending as IETF Proposed Standard RFC.
19. J. Korhonen, and T. Savolainen: EDNS0 Option for Indicating AAAA Record Synthesis and Format. IETF Internet Draft, 2011, work-in-progress.
20. IETF Behavior Engineering for Hindrance Avoidance (BEHAVE) Working Group: `http://tools.ietf.org/wg/behave/`
21. Trilogy Project: `http://www.trilogy-project.org/`
22. Hub-of-All-Things (HAT) Project: `http://hubofallthings.wordpress.com/`
23. Jon Crowcroft: Presentation on HAT Project. Internet Science Conference 2013. `http://www.youtube.com/watch?v=9Zq3zHLptpI`