

Bridging the Gap between Internet Standardization and Networking Research

Aaron Yi Ding^{†**}, Jouni Korhonen[‡], Teemu Savolainen[§], Markku Kojo^{*}, Jörg Ott[◊]
Sasu Tarkoma^{**}, Jon Crowcroft[†]

[†]University of Cambridge

^{*}University of Helsinki

[‡]Broadcom

[§]Nokia

[◊]Aalto University

[•]Helsinki Institute for Information Technology HIIT

{aaron.ding, jon.crowcroft}@cl.cam.ac.uk, {kojo, sasu.tarkoma}@cs.helsinki.fi, jorg.ott@aalto.fi
jouni.korhonen@broadcom.com, teemu.savolainen@nokia.com

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.

The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

The participation of the network research community in the Internet Standards Development Organizations (SDOs) has been relatively low over the recent years, and this has drawn attention from both academics and industry due to its possible negative impact. The reasons for this gap are complex and extend beyond the purely technical. In this editorial we share our views on this challenge, based on the experience we have obtained from joint projects with universities and companies. We highlight the lessons learned, covering both successful and under-performing cases, and suggest viable approaches to bridge the gap between networking research and Internet standardization, aiming to promote and maximize the outcome of such collaborative endeavours.

Categories and Subject Descriptors

C.2.6 [Computer-Communication Networks]: Internet-working—Standards

General Terms

Documentation; Standardization

Keywords

Internet Standardization; Networking Research

1. INTRODUCTION

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

To show the existing gap between the research community and the Internet SDOs, we first analyse the participation ratio of researchers in one of the largest Internet SDOs, the Internet Engineering Task Force (IETF), which used to have a

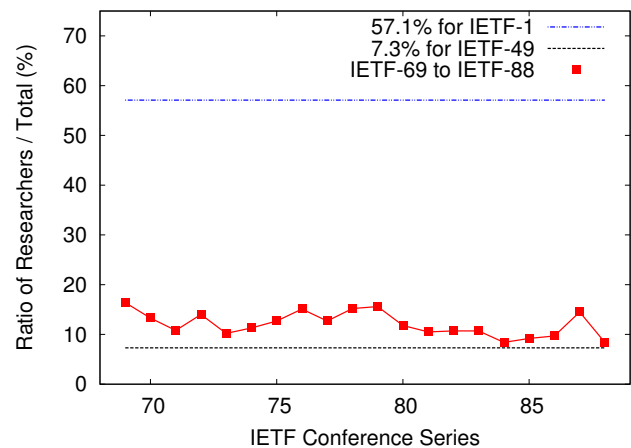


Figure 1: Researcher Participation at the IETF

high percentage of the researcher participation.¹ As shown in Figure 1 covering the most recent 20 IETF meetings, the ratio remains relatively low, ranging from 8.4% to 16.4%.² Being the most popular meeting in the IETF history with 2810 attendees, the IETF-49 yields only 7.3% even by taking into account all the professionals from the academia and industrial labs. When comparing to the first IETF in 1986, where the ratio is 57.1%, the current situation is far from encouraging.³

As a guidance for protocol design, RFC 5218 [1] summarizes a set of key factors and requirements for successful global deployment. However, except for general discussions on the transfer of R&D results to commercial products in IT business, there are few dedicated studies or investigations on how to bridge the existing gap between networking

¹The term researcher in this paper refers to the professionals who work mainly in the academia.

²The ratio includes also the professionals from the research institutes of industrial companies.

³Note that the Internet was commercialized in 1990s, resulting in a significant increase in the non-researcher participation at the IETF.

research and Internet standardization. In seeking feasible solutions to bridge such a 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 professionals participating in our joint projects, we have achieved good results in terms of academic publications and standardization contributions. 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 challenges and opportunities for the collaboration between networking research and Internet SDOs (mainly on the 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 collaborations [2, 3, 4, 5]. Our case study focuses on two specific topics, mobile traffic offloading and IPv6 transition technologies, which recently have drawn great attention from both research communities and SDOs.
3. We further offer our insights on the lessons learned from the successful and under-performing 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 discuss the latest trend in Section 6. We conclude our paper in Section 7.

2. CHALLENGES AND OPPORTUNITIES

To enhance the technology transfer in IT industry, there are prior works investigating how to promote the collaboration for commercial success [6, 7, 8]. However, we are facing unique challenges for the existing gap between networking research communities and Internet SDOs that require a dedicated analysis. Based on our experience through joint projects, we identify the main challenges including technical and non-technical aspects and discuss the opportunities.

2.1 Technical Challenges

From the technical viewpoint, we observe four visible challenges:

- **Process and Target:** The scientific publishing and standards development follow different processes. The first visible difference is the review process adopted by the research community and SDOs such as the 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. For IETF standardization, it can take several years for a proposal to be approved as an RFC standard. The typical 3-5 years PhD cycle also contributes to the disconnect since the student may have graduated already till the time when the standardization process approaches the final stage. 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.

- **Contribution Control:** Different from research publications, where authors have the right to decide the core content, standardization work is often conducted through working groups based on consensus and is hence largely controlled by the standardization community rather than the authors. This may demotivate researchers for participation and prolong the progress.
- **Methodology:** Rather than fixing the specific problems in the Internet, researchers tend to prefer re-inventing the Internet, redefining the architecture or conducting pure analytic evaluation with no protocol impact. Therefore, many popular research topics do not fit to the agenda on which the Internet SDOs focus.
- **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. 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.

2.2 Non-technical Challenges

Concerning the non-technical aspect, we observe several challenges that are more subtle and harder to quantify. These include:

- **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 professionals face similar challenges for promotion when they mainly publish scientific articles.
- **Performance Evaluation:** There is a clear gap between the performance metrics of evaluation used in academia and industry. For academic researchers, the major metrics include scientific publications, citations, 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 patents, 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 focus for standardization professionals.
- **Financial Support:** Lack of financial support and understanding from direct superiors or institute leaders is another inevitable challenge that directly affects researchers' motivation to attend SDO meetings to learn from and collaborate with other professionals.

Table 1: Project Outcome for Mobile Traffic Offloading

	Research	Standardization
Contributions	5 publications in good venues	1 IETF Internet draft, 1 Change Request in 3GPP SA2 (System Aspects Working Group 2)
Performance	Very good	Under-performing

- Fear: There is a subtle overestimation for the difficulty of the work in SDOs, 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 statements either in conference seminars or standardization meetings that the work from the research community is unrealistic and incomprehensible, while industrial work being purely money-driven and short sighted.

2.3 Opportunities

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 contributions 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 community.

Concerning joint benefits, researchers often need testing facilities and real world problems that can be supported by industrial partners. For example, companies can provide access to core network facilities, usage traces, and challenging problems that are regarded as rare and valuable input for the academic research. Obtaining research input from SDOs such as IETF can also be very helpful. Meanwhile, companies focusing on standardization need strong analytical tools and expertise from the 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 outcomes 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 professionals and academic researchers to work together and learn from each other.

One opportunity to highlight is the personnel transfer, which often covers 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 the size of problem. There are also intermediate ways such as sending students to industry through internship. Another good example, as a motivation, is the job opportunity created for PhD students who actively participate in the IETF work and hence obtain valuable experience for their future ca-

reer. The recent trend for industrial professionals moving back to academia also shows the potential for sharing expertise and experience. For industry, the PhD hires at top Internet companies often represent good investment.

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 2012, we investigated this domain from both the research and the standardization point of view.

For research, we achieved good results [9, 10, 11, 12, 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 the 3GPP standards development organization. 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 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 approaches, one based on DHCPv6 and the other based on IPv6 Neighbour 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 many IETF participants or even 3GPP delegates. We summarize the outcome in Table 1.

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 *Good Technical Design* (RFC 5218 Section 2.1.7.) since they, for example, mix IPv4 into IPv6 only protocols and introduce routing style functionality into the IPv6 Neighbour Discovery protocol. Both are typically hard to justify in the IETF. When it comes to the 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,

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

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

but the standardization effort can be regarded as an under-performing case. Frankly, the authors were aware that the offloading topic is challenging to drive in standardization, specifically in the IETF working groups that are not already familiar with cellular network technologies. Almost the same could be said about 3GPP. Proposing a solution that does not fit to an existing mindset is usually a failure to begin with, unless there is a really strong hype behind 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 a 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 an Informational RFC. Based on the community feedback and our feasibility test, the heuristic discovery approach has been adopted in the IETF transition toolbox for the IPv6 Internet as a Standards Track RFC. We highlight the outcome in Table 2.

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 for its impact on other entities in the network. Although the heuristic discovery offers only moderate adap-

tation, 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.

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

5. LESSONS AND SUGGESTIONS

5.1 Lessons from Projects

Through our collaborative projects investigating mobile traffic offloading and IPv6 transition technologies, we have learned several lessons. For the under-performing case on mobile traffic offloading, we list the key findings:

- The offloading problem is well comprehended in the 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 professionals are subjected to various restrictions such as time commitment and contract duration. On the other hand 3GPP had already selected their preferred *telecom style* solutions to the offloading topic and trying to convince them to revert to something that basically needs no new work or equipment was not welcome. Because offloading is essentially a system topic, which also creates confusion in the IETF due to the mismatch of the scope. Furthermore, in our case, the authors' priority and interest gradually shifted 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 the work is meaningful and suitable as proposed standards. We chose the IETF Multiple Interface (MIF) working group [15] as a very challenging place 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 neighbour 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 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.

5.2 Suggestions

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 the thinking, simulation, and prototyping but took very active role also in standardization. With the support from EU partners and companies, the proposal is gradually integrated into commercial OS versions such as iOS. A key to such success is the patience and investment in time, which can span very long periods of time (5 to 10 years). Meanwhile, comparing to the large projects, smaller projects offer the advantage that 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].
2. Identifying the Targets: Many research proposals presented at the IETF received negative feedback mainly because the work items do not fit in to the working group agenda nor the standardization scope. It is important to identify which part of the research work can be carried forward as standardization. This requires time and effort to investigate the scope of targeted SDO and its standardization process. It is highly recommended, if possible, to attend the SDO meetings in person to have the opportunity to exchange opinions directly with standardization professionals and consult them for advice.
3. Evaluating the Impact: Being a classical factor to evaluate the research impact, citation suffers from the limitation that it mainly reflects the influence received by the peers in the academic community. At the same time, the standardization work, even the ones including significant research components, is given little value, despite that technical standards may receive more reviews than typical academic papers. As one of the barriers for researchers to participate in the standardization work, a more comprehensive evaluation metrics can motivate researchers to actively get involved, for instance, by taking the standards contribution and its industrial impact into account. One step forward in promoting the value and role of the Internet RFCs as scholarly publications is discussed in [23].
4. Understanding the Difference: If a problem is well defined, being real in existing networks but relatively easy to fix, the standardization goes smoothly. However, if a problem is hard and the fixing process is complex, the authors would have to prepare much better problem descriptions, do much more footwork in convincing people of the problem and the feasibility of a solution, and generally more ready to spend time in selling the solution. For instance, to get such a solution through the 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 where the problem itself is more valuable than the complexity of deployment and solutions. Another difference is the time commitment. While research papers can take 1+ years to appear in a conference or a journal, in an SDO, publishing an idea is very fast, but making the idea as a standard can take much longer. In our IPv6 heuristic discovery proposal [18], it took from September 2010 to October 2013, with 20 revisions leading to the final RFC, i.e., 21 document versions and *three years* in total. Such a time frame can be too long for a researcher to be involved, given the fact that 1-2 year project funding is common nowadays, which may demotivate the researchers even to start the process.
5. Organizational Support: For industry professionals, 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 find little value to contribute to 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 professionals 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 are 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 chal-

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

lence of career promotion for both researchers and industrial professionals.

6. DISCUSSIONS AND OUTLOOK

One myth about Internet SDOs such as the IETF is the time requirement to publish a standard document. For comprehensiveness, we analyse the trend of the recent 100 IETF RFCs (RFC 6931 to RFC 7038) to highlight how long it often takes starting from the initial Internet Draft to the final document. As illustrated in Figure 2, around 50% of RFCs we investigated demand more than two years minimum devotion, where the longest one goes even up to 137 months, i.e., 11+ years! Among our samples, only 20% demand less than a year with the minimum of 5 months requirement (most of them are not technical standards). Comparing to the typical 6-month to 1 year conference publication agenda in the academic domain, it is clearly a time-consuming journey to publish Internet standards. On top of that, there are also an intense initial period before the document process and various invisible elements, including the time for selecting or setting up working groups, initiating the topics in the community, and day-to-day time commitment to keep up with the relevant discussions in the working groups. One good sign we observe is that among the recent 100 RFCs, researchers made contributions to 36% of them.

Meanwhile, one interesting and surprising finding is on the publication openness where many Internet SDOs have been making documents openly accessible, but in the networking research domain, the open-access procedure starts only recently (except for SIGCOMM and CCR which have

published the accepted papers online for many years without the subscription requirement). On the other hand, 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. In particular for academic publications, a visible trend is that many journals start open calls for industrial contributions in terms of standards and deployment experience. To highlight the linkage between standardization and research, we summarize our observations (non-exclusive) in Table 4.

The Internet Research Task Force (IRTF) [24] is a organization parallel to IETF that produces Informational and Experimental documents [25]. For many years, IRTF has been focusing on long term research issues related to the Internet and striving to attract professionals from both academia and industry. Although the full potential of IRTF is yet to be exploited, since IETF has moved from a 'researcher-friendly' forum to an industry-dominant standardization unit, IRTF could bridge the gap by providing a channel for researchers to obtain input and feedback from industry and transfer their research ideas to standardization.

One activity we want to highlight is the recent diversity debate at the IETF. As initiated and supported by the IETF Chair, the IETF has set up a design team to improve the participation diversity for maintaining high standards of quality in the long run [26]. The agenda and work items are now going through discussions among researchers and industrial professionals [27]. One of the main goals is to bring researchers into the IETF for standardization contributions by seeking ways to create exchange opportunities between industry and academia, e.g., through the IETF University Outreach Pilot Programme [28], the ISOC Fellowship program [29], and the Applied Networking Research Prize [30]. Although the final result is not clear yet, we view it certainly as moving toward a promising direction.

7. CONCLUSION

This editorial presents our exploration of how to bridge Internet standardization and networking research. We believe there is great potential for comprehensive collaboration between the research and SDO communities and hope our lessons and suggestions can shed light on this challenging domain and further promote the sustainable development of the Internet.

Acknowledgments

This work has been carried in the WiBrA project funded in part by the TEKES Finland, in collaboration with Nokia, NSN, and TeliaSonera. The authors are grateful for the comments from Jukka Manner, Dina Papagiannaki, Colin

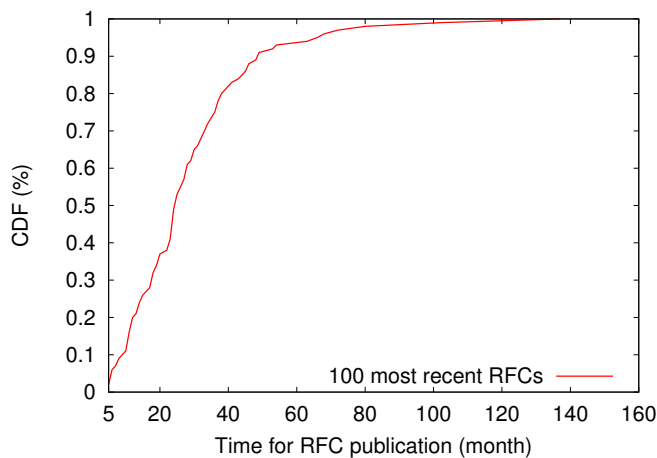


Figure 2: Time Required for RFC Publication

Perkins, Pasi Sarolahti, Henning Schulzrinne, Joe Touch, and Michael Welzl. The PhD research of Aaron Yi Ding is supported by the Academy of Finland, HIIT, and Nokia Foundation.

8. 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] 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] 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. Internet RFCs, ISSN 2070-1721, RFC 7051, 2013.
- [18] T. Savolainen, J. Korhonen, and D. Wing: Discovery of IPv6 Prefix Used for IPv6 Address Synthesis. Internet RFCs, ISSN 2070-1721, RFC 7050, 2013.
- [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] Brian E. Carpenter, and Craig Partridge: Internet Requests for Comments (RFCs) as Scholarly Publications. *ACM SIGCOMM Computer Communication Review*, 40(1):31–33, January 2010.
- [24] Internet Research Task Force (IRTF): <http://irtf.org/>
- [25] A. Falk: Definition of an Internet Research Task Force (IRTF) Document Stream. Internet RFCs, ISSN 2070-1721, RFC 5743, 2009.
- [26] IETF Design Team Call for Community Input: <http://www.ietf.org/mail-archive/web/ietf/current/msg80076.html>
- [27] Diversity Design Team Wiki: <https://wiki.tools.ietf.org/group/diversity-dt/wiki>
- [28] IETF University Outreach Pilot Programme: <http://www.internetsociety.org/doc/ietf-university-outreach-briefing-paper>
- [29] Internet Society Fellowship to the Internet Engineering Task Force (IETF) Programme: <http://www.internetsociety.org/what-we-do/education-and-leadership-programmes/ietf-and-ois-programmes/internet-society-fellowship>
- [30] Applied Networking Research Prize: <http://irtf.org/anrp>