

Report on the Workshop on Research and Applications of Internet Measurements (RAIM)

Mat Ford
Internet Society
ford@isoc.org

Lars Eggert
NetApp
lars@netapp.com

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The author takes full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

This report is a brief summary of the Internet Research Task Force and Internet Society workshop on Research and Applications of Internet Measurements (RAIM) in co-operation with ACM SIGCOMM that took place on Saturday, October 31, 2015 in Yokohama, Japan. The workshop provided an opportunity for researchers and practitioners in the field of Internet measurements to become acquainted and share their work.

1. INTRODUCTION

The RAIM workshop had its genesis in the mutually reinforcing desires of the ACM SIGCOMM Executive Committee to foster more industry participation in the SIG and of the Internet Engineering Task Force (IETF) and the Internet Research Task Force (IRTF) to foster more academic participation in the IETF and IRTF. Desired outcomes for the workshop included information exchange and fostering more continuous interaction and collaboration between research, standards and engineering.

Although RAIM was an invitation-only workshop for logistic reasons, a stated goal of the organizers was to make it especially easy for researchers to attend and present. For this reason, they were asked to submit a reference to a relevant, recently-published academic paper, with an option to submit an original short position paper. The agenda was also deliberately structured to include a lot of time for interactive discussion.

The workshop was structured into three sections: measurement platforms & tools, measurement of fixed & mobile broadband access networks, and Internet characterization through measurements. Each section began with a session keynote followed by a series of short, three-minute lightning talks that allowed everyone who had a contribution to the topic a chance to speak and to keep the session very dynamic. A lengthy discussion period on the session topic followed all the lightning talks.

The RAIM workshop had 72 registered attendees, roughly a third of which were industry participants. In addition to the three keynote talks, a total of 44 lightning talks were given. All accepted papers, together with all presented slides and full recordings of the workshop are available on the RAIM workshop web page [1].

Although the RAIM workshop was likely a one-off event that was enabled by the proximity in time and space of the ACM IMC 2015 conference in Tokyo and the IETF-94 meeting in Yokohama, the Internet Society, IETF and IRTF offer a number of other ways for academics to get involved in applied Internet engineering. They include Internet Society fellowships [2], the Applied Networking Research Prize (ANRP) [3], as well as a new annual Applied

Networking Research Workshop (ANRW) series currently under preparation [4].

2. ABOUT THE IETF AND IRTF

This section gives a brief overview about the IETF and IRTF, in order to provide some context about the workshop. Readers familiar with these organizations may want to skip ahead to Section 3.

As described on the IETF web page [5], the IETF “is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. It is open to any interested individual. The actual technical work of the IETF is done in its working groups, which are organized by topic into several areas (e.g., routing, transport, security, etc.). Much of the work is handled via mailing lists. The IETF holds meetings three times per year.” Meetings attract typically around 1000-1500 attendees, with many more participating in IETF standardization by email. A longer description of the IETF [6] and information for newcomers [7] are available.

The IRTF is a sister organization to the IETF and acts as its research arm. It has around a dozen research groups that focus on longer term research issues related to different aspects of the Internet. Unlike IETF working groups, IRTF research groups do not publish standards, and therefore need not have consensus on the results they publish.

Both the IETF and the IRTF publish technical and organizational documents about the Internet, including the specifications and policy documents, in the RFC series, alongside RFCs by the Internet Architecture Board (IAB) and independent submissions by the wider community. The Internet Assigned Numbers Authority (IANA) acts as the central coordinator for the assignment of unique parameter values for IETF and IRTF protocols.

3. MEASUREMENT PLATFORMS & TOOLS

The first session of the RAIM workshop identified and discussed the various platforms and tools available to researchers and practitioners interested in Internet measurements.

3.1 SESSION KEYNOTE

Brian Trammel (ETHZ) gave the session keynote and talked about the different dimensions of work in tools and measurement platforms, illustrated the state-of-the-art by reference to the mPlane

project, and provided some comments about likely future paths forward.

Measurements can provide different views of the Internet, e.g. routing topology, addressing and naming, security, dataplane performance or impairment, and traffic matrices. There are also different reasons to conduct measurements depending on whether one is working in network operations, research or engineering.

More work appears to be needed in the areas of techniques and methodologies for measurement, representing and analyzing the results, and coordination of multiple measurements. Brian introduced mPlane as an example of a platform that uses an error-tolerant RPC protocol connecting clients with components to cooperatively perform network measurements and analysis using heterogeneous tools. RIPE ATLAS was also introduced as an operationally focused active measurement platform using thousands of distributed probes. Observatories to centralize partially processed data for a specific purpose were identified as one way to address some of the challenges of centralization and data-sharing. Better instrumentation and providing out-of-band access to statistics will grow in importance. Explicit exposure of information as part of every protocol exchange could yield lots of valuable data with very low sampling rates.

In conclusion, Brian identified measurement platforms as providing the engineering to fill the gap between research and practice both in terms of coordination to allow measurement to scale and representation yielding portable and universally applicable results. He urged us to move towards treating measurement as a first-class function of the stack.

3.2 LIGHTNING TALKS

The measurement platforms & tools session featured seventeen lightning talks. Al Morton (ATT) talked about the interactions between measurement and the current trends of network function virtualization and software-defined networking. Mirja Kühlewind (ETHZ) introduced the goals and requirements for a middlebox observatory. Anna Maria Mandalari (UC3M) discussed using crowdsourcing platforms to conduct large-scale Internet measurement campaigns. Jeferson C. Nobre (UFRGS) talked about using traffic information to improve the detection of SLA violations through a P2P approach. Joachim Fabini (TU Wien) presented the Representative Delay Measurements Tool intended to detect and correct for timing impairments in delay measurements. Andra Lutu (Simula) introduced the MONROE platform for measurements and experiments in operational mobile broadband networks in Europe. Alessandro Improta (IIT) introduced the Isolario project intended to motivate sharing BGP data to obtain real-time services. Paul Emmerich (TU Munich) presented MoonGen, a scriptable high-speed packet generator intended to combine the advantages of software (cheap, flexible) and hardware (precise) packet generators. Matthew Luckie (U Waikato) presented the scamper packet prober tool. Robert Kisteleki (RIPE-NCC) introduced the RIPE ATLAS measurement platform's latest statistics and new features. Liam McNamara (SICS) presented the CheesePi measurement project that aims to quantify and characterize the Swedish Internet in collaboration with the regulator in that country. Thomas Holterbach (ETHZ) talked about interference between measurements on the RIPE ATLAS platform. Kazunori Fujiwara (JPRS) presented some of the datasets available for DNS data collection and analysis and some results obtained by combining these sources. Hirochika Asai (U Tokyo) talked briefly about the operational monitoring and measurement challenges of large-scale WiFi networks. Ramakrishna

Padmanabhan (U Maryland) presented some interesting observations about round-trip time (RTT) latency measurements. Steve Rich (Cisco) presented some of the funding opportunities available from Cisco's Advanced Security Research Group. Finally, Rolf Winter (HS Augsburg) described his research into the data exposure risks of multicast and broadcast traffic.

3.3 DISCUSSION

The follow-on discussion ranged across the apparent duplication of effort in some of the presented measurement platforms and the need for there to be a healthy ecosystem of measurement tools and platforms to avoid the risk of a single standard being gamed. The ethics of some of the studies described was also touched upon and this was identified as an area where new norms are rapidly evolving. A clear need was expressed for recording metadata regarding both measurements and tools. Global searchability across multiple repositories of measurement results was identified as highly desirable. The differential behavior of some networks based on different types of traffic was noted and the interaction of this aspect with the increased use of encryption was also raised as an important development for measurement tools to address.

4. MEASUREMENT OF FIXED & MOBILE BROADBAND ACCESS NETWORKS

The second session of the workshop addressed the state-of-the-art and challenges relating to the measurement of fixed and mobile broadband access networks.

4.1 SESSION KEYNOTE

In his session keynote talk, Phil Eardley (BT) spoke about Quality of Experience (QoE) and noted that QoE is about much more than just raw throughput. Phil presented an architecture for measuring QoE using both active and passive measurements. Large scale active measurements are helping BT handle core broadband traffic growth of 65% year-on-year by identifying hotspots in the network and understanding the impact and operation of new devices, technology, products and services. Other ISP use cases include identifying and isolating network faults, monitoring suppliers and understanding the customer's end-to-end service experience. Regulators and end-users also have their own use cases for the measurement tools.

Phil illustrated the use of measurements to help diagnose issues with the BBC iPlayer catch-up TV service and the use of in-network caching. Improving measurement of QoE will require more realistic tests for video and voice, and ensuring metadata is accurate. Conducting tests on a per-line basis has potential benefits too but finer granularity means more probes are required and is motivating a shift from hardware to software probes. Improved data analysis including on-demand testing, improved diagnostics and standardization were all identified as areas for further research and development.

Phil concluded his talk by presenting the standardization landscape for broadband measurement, including the tests standardized in the IETF IP Performance Metrics (IPPM) working group, the protocols standardized in the IETF Large-Scale Measurement of Broadband Performance (LMAP) working group and their relationship with the work of the Broadband Forum.

In the discussion following the keynote presentation, the challenges of capacity testing using packet pairs were discussed and the potential for hybrid measurements as a way to tackle them was identified. The importance of home network testing was also supported as a hard problem to address. For supply chain mapping,

network tomography was identified as a tool developed by the research community that could be applicable in operational settings. Finally, it was agreed that aligning incentives to encourage collaborative measurements is key.

4.2 LIGHTNING TALKS

Nine lightning talks were presented on this topic area. Joachim Fabini (TU Wien) presented access network measurements that show that real access networks do not always conform to the models often used by researchers. Renata Teixeira (INRIA) presented work to develop easy-to-use home network diagnosis tools, including passive measurements at Internet gateways and browser-based collaborative measurement tools. Matthew Luckie (U Waikato) briefly talked about a joint project between CAIDA and MIT investigating the nature and prevalence of interdomain congestion. Anna Brunstrom (U Karlstad) presented measurements from four Swedish mobile operators showing the impact of background TCP flows and the much greater sensitivity of short flow completion times to changes in RTT as compared with changes in throughput. Deliang Chang (Tsinghua) talked about a study using DNS logs to perform OS fingerprinting and observe the prevalence of NAT/tethering usage in a large-scale network. Leslie Daigle (Thinking Cat) emphasized the importance of collaboration between organizations and collation of data to manage privacy concerns as ways to measure the Internet's stability and health. Pedro Casas (FTW) presented a platform to measure QoE in cellular networks using a combination of passive DPI probes and distributed active measurements. Rocky Chang (Hong Kong Poly) briefly presented three studies that help to highlight some of the challenges of measuring networks accurately. Kensuke Fukuda (NII/Sokendai) described research tracking the growing use of WiFi offload by mobile smartphone users and the impact of this trend on residential broadband in Japan.

4.3 DISCUSSION

In group discussion following the lightning talks, the challenges of meaningfully measuring QoE were discussed at some length. The extreme heterogeneity of end user environments was identified as a huge challenge for extrapolating meaningfully from QoE research to operational settings. The various mobile network studies also showed clearly that assuming constant bitrate links and FIFO queues was increasingly likely to yield erroneous results.

Data sharing was a topic identified by the group for discussion. It was noted that funding agencies are increasingly bringing pressure to bear on researchers to improve data sharing during and after the life of research projects, although it was also noted that regulatory and privacy concerns inevitably complicate matters. Lack of common formats and indexed metadata were identified as problems and the ability to review datasets without having to formulate a research question in advance would help. Industry could help more by sharing data and contributing some 'grand challenge' type questions.

5. INTERNET CHARACTERIZATION THROUGH MEASUREMENTS

The final session of the workshop discussed the latest work to characterize the Internet at scale through measurements.

5.1 SESSION KEYNOTE

In the final session keynote of the workshop, Ethan Katz-Bassett (USC) talked about ways to drive up the impact of Internet measurement studies. It is not enough to place some measurement

probes on the network, gather a dataset, pose some research questions, analyze the data, and publish a paper. To have real impact the measurement results need to influence Internet operations.

Ethan cited the example of TCP Gentle Aggression to support his argument. This work, which resulted in a reduction in mean response times of 23% for Google and upstreamed a new default mechanism in Linux 3.10+, benefited from appropriate data and vantage points, the contribution of a real-world problem and a clear pathway to deployment from Google. Most measurement research is challenged in all these areas: there are limited options for evaluation, limited data shapes the analysis, and there is limited visibility into real-world operational concerns. One way to increase impact by tackling some of the difficulties in performing measurements themselves is to provide long-running measurements, tools and testbeds. Mapping Google's expansion and the PEERING BGP testbed were presented as examples.

Supporting long-running measurement services is not without its own set of challenges and costs and some ways of encouraging the community to provide more and better support for long-running services were discussed. Ethan noted that IMC already mentions 'advances in...facilitating [data] sharing' in the CFP, has a community session for informal advertisements of new datasets and tools, and awards the best new dataset made publicly available.

Some ideas for new ways to encourage data, code and testbeds to be made public were also discussed. The ability of IETF work in the form of Internet-Drafts and RFCs to influence operations, provide support for testbeds and tools, and to inspire research questions was noted. Finally, some suggestions for new ways that the IETF could help to support impactful Internet measurement research were made, including publishing a list of important measurement questions, sharing data, hosting vantage points, and encouraging communication and collaboration between communities (through joint conferences, Applied Networking Research Prize, student scholarships to IETF, etc.).

In the discussion after Ethan's keynote, the challenges for academic researchers to justify attending IETF meetings, and knowing how to navigate them once there were discussed at some length. The ability for IRTF research groups to operate differently, and to hold meetings outside IETF meetings was noted as a way to address some of these obstacles to more involvement. The proposed Measurement and Analysis for Protocols (MAP) research group was identified as a potential 'landing pad' for Internet measurement work in the context of the IETF.

5.2 LIGHTNING TALKS

The lightning talk portion of the session consisted of eighteen talks. Nevil Brownlee (U Auckland) presented the results of research to characterize the performance and fairness of big data transfer protocols on long-haul networks. Qinwen Hu (U Auckland) talked about some of the implications of interface identifier allocation mechanisms in IPv6. Emile Aben (RIPE-NCC) briefly discussed a framework to deliver mesh-traceroutes for a selection of RIPE Atlas probes in a set of countries and to provide automated analysis. Giovane C. M. Moura (SIDN) presented some initial results of a study modelling and measuring Internet-wide DHCP churn. Matthew Luckie (U Waikato) presented the results of investigations into TCP behavior using tbit tests built into the scamper tool and solicited feedback from the IETF on new tests to pursue in the coming year. Colin Perkins (U Glasgow) motivated using ECN with UDP-based protocols and presented some results from an initial measurement study that showed evidence of middlebox

interference on ~1% of paths. Zeqi Lai (Tsinghua) talked about the efficiency of the protocols used for synchronization in popular cloud storage services and motivated future work on a standardized sync protocol. Maria Ines Robles (Ericsson) presented some considerations regarding device management and measurements in the Internet-of-Things. Ruwaifa Anwar (Stony Brook) discussed the results of work to analyze the validity of the Gao-Rexford model of interdomain routing using RIPE Atlas probes to measure the applicability of the model to today's Internet. Megumi Ninomiya (IIJ) presented work to model and measure the playback delay for HTTP-based live video streaming. Christoph Paasch (Apple) presented some lessons learned from studies of TCP Fast Open and ECN that showed the importance of measuring the whole flow and not just the handshake to observe middlebox interference. Yves Vanaubel (U Liège) presented the Label Pattern Recognition algorithm, a method to analyze traceroute data including MPLS information to better understand the transit path diversity deployed within a given ISP. Alberto Dainotti (CAIDA) presented results of a detailed and ongoing study into IPv4 address space utilization. Satoshi Kamei (NTT Communications) talked about the results of an active measurement study that helped to shed light on the use of different content delivery platforms in major content delivery events. Kevin Fall (CMU) presented some thoughts about Compact Routing (CR) and whether a CR protocol could be developed for Internet use. Eric Osterweil (Verisign) talked briefly about work to model the systemic dependencies between WebPKI and DANE approaches to web security using attack surface analysis. Stephen Strowes (Yahoo!) talked about some of the disparity in IPv6 uptake between countries, between networks, and between fixed and mobile deployments. Varun Singh (callstats.io) presented the results of detailed analysis of WebRTC measurements in real-world deployments.

5.3 DISCUSSION

In the group discussion following the lightning talks, it was observed that the community spends most of its time measuring applications and services considered 'too big to fail', e.g., DNS or the web. Higher-layer or more niche applications are considered less important and therefore less interesting. However, niche applications are subject to research studies when they get very big, e.g., Netflix. There was interest and discussion of the research results from CAIDA showing that a substantial portion of routed IPv4 address space is unused. Sharing more data (without revealing sensitive information) was identified as the best way to keep refining this result.

The motivations for and timeliness of standardization were discussed and it was observed that standardization typically

happens when people want to interoperate, and the IETF likes it when people who would deploy a standard are involved.

6. CONCLUSION

In concluding the meeting, Lars Eggert noted that although the RAIM workshop was most likely to be a one-off event taking advantage of the proximity of the ACM IMC conference in Tokyo and the IETF meeting in Yokohama, he was still very interested in increasing the intersection between the communities identified in the introduction to the workshop. The MAP RG mailing list was identified as the best venue for continuing the discussion between operators and researchers [8].

ACKNOWLEDGMENTS

In addition to the authors, the RAIM 2015 workshop was organized by Mark Allman (ICIR), Kenjiro Cho (IIJ), Ratul Mahajan (Microsoft Research), Renata Teixeira (INRIA), Brian Trammell (ETHZ) and Darryl Veitch (UTS). Funding for the workshop was provided by ACM SIGCOMM, the Internet Society and the Internet Research Task Force (IRTF).

Lars Eggert has received funding from the European Union's Horizon 2020 research and innovation program 2014-2018 under grant agreement No. 644866 ("SSICLOPS"). This paper reflects only the authors' views and the European Commission is not responsible for any use that may be made of the information it contains.

REFERENCES

- [1] RAIM Workshop Web Site. <https://irtf.org/raim-2015>
- [2] The Internet Society Fellowship to the Internet Engineering Task Force (IETF). <http://www.internetsociety.org/what-we-do/education-and-leadership-programmes/ietf-and-ois-programmes/internet-society-fellowship>
- [3] Applied Networking Research Prize. <https://irtf.org/anrp>
- [4] ACM, IRTF & ISOC Applied Networking Research Workshops. <https://irtf.org/anrw/>
- [5] About the IETF. <https://www.ietf.org/about/>
- [6] The Tao of IETF: A Novice's Guide to the Internet Engineering Task Force. <https://www.ietf.org/tao.html>
- [7] Getting Started in the IETF. <https://www.ietf.org/newcomers.html>
- [8] MAP RG Mailing List. <https://www.irtf.org/mailman/listinfo/maprg>