

# Global Measurements: Practice and Experience

Edited by

Vaibhav Bajpai<sup>1</sup>, Arthur W. Berger<sup>2</sup>, Philip Eardley<sup>3</sup>, Jörg Ott<sup>4</sup>,  
and Jürgen Schönwälder<sup>5</sup>

1 Jacobs University – Bremen, DE, v.bajpai@jacobs-university.de

2 Akamai Technologies – Cambridge, US, awberger@csail.mit.edu

3 BT Research – Ipswich, GB, philip.eardley@bt.com

4 TU München, DE, ott@in.tum.de

5 Jacobs University – Bremen, DE, j.schoenwaelder@jacobs-university.de

---

## Abstract

This article summarises a 2.5 day long Dagstuhl Seminar on “Global Measurements: Practice and Experience” held in January 2016. This seminar was a followup of the seminar on “Global Measurement Frameworks” held in 2013, which focused on the development of global Internet measurement platforms and associated metrics. The second seminar aimed at discussing the practical experience gained with building these global Internet measurement platforms. It brought together people who are actively involved in the design and maintenance of global Internet measurement platforms and who do research on the data delivered by such platforms. Researchers in this seminar have used data derived from global Internet measurement platforms in order to manage networks or services or as input for regulatory decisions. The entire set of presentations delivered during the seminar is made publicly available at <http://materials.dagstuhl.de/index.php?semnr=16012>.

**Seminar** January 4–7, 2016 – <http://www.dagstuhl.de/16012>

**1998 ACM Subject Classification** C.2.3 Network Operations, C.4 Performance of Systems

**Keywords and phrases** Internet measurements, Quality of experience, Traffic engineering

**Digital Object Identifier** 10.4230/DagRep.6.1.15

## 1 Executive summary


*Vaibhav Bajpai*

*Arthur W. Berger*

*Philip Eardley*

*Jörg Ott*

*Jürgen Schönwälder*

**License**  Creative Commons BY 3.0 Unported license

© Vaibhav Bajpai, Arthur W. Berger, Philip Eardley, Jörg Ott, Jürgen Schönwälder

Several large-scale Internet measurement platforms have been deployed during the last years in order to understand how the Internet is performing, to observe how it is evolving, and to determine where failures or degradations occur. Examples are the CAIDA Archipelago (Ark) platform [6] (used for Internet topology discovery and detecting congestion on interdomain links), the SamKnows platform [4] (used by regulators and network operators to study network performance), the RIPE Atlas platform [3, 5] (that provides measurement services to network operators and researchers), the Netradar system [8] (for performing wireless performance measurements), and the BISmark project [9]. European collaborative research



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 3.0 Unported license

Global Measurements: Practice and Experience, *Dagstuhl Reports*, Vol. 6, Issue 1, pp. 15–33

Editors: Vaibhav Bajpai, Arthur W. Berger, Philip Eardley, Jörg Ott, and Jürgen Schönwälder



DAGSTUHL  
REPORTS

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

projects lately have been working on a Measurement Plane (mPlane) [10] and how to incorporate measurement results into network management systems (e.g., Leone) [2]. Related projects (e.g., Flamingo) [1] are increasingly working with measurement data from these platforms. Large-scale measurements are meanwhile also used to drive network operations or to dynamically adjust how services are delivered to customers. Content Delivery Network (CDN) providers use measurement data to optimize content caches and to tune load balancing algorithms. One key challenge is that global Internet measurement systems can generate large amounts of data that need to be processed to derive relevant information.

This seminar (#16012) was a followup of the Dagstuhl seminar on Global Measurement Frameworks (#13472) [7]. The main focus of the first seminar was an exchange of ideas on the development of global measurement infrastructures, frameworks and associated metrics. Some of this work is now further pursued in standardization bodies [4] such as the IETF Large-Scale Measurement of Broadband Performance (LMAP) working group and the Broadband Forum. The goal of this followup seminar was to focus on the experience obtained with different metrics, tools, and data analysis techniques. It provided a forum for researchers to exchange their experience with different practices to conduct global measurements. The aim was to identify what works well in certain contexts, what has proven problematic in other contexts, and identify open issues that need further research. The seminar approached this by looking at three distinct dimensions: (a) Measurement metrics, (b) data processing technologies and (c) data analysis methodologies. Some key questions were:

1. Which metrics have been found useful for measuring Quality of Experience (QoE) of certain classes of services? Which metrics have been found problematic? Is it possible to find indicators for good metrics and problematic metrics?
2. Which technologies have been found useful for storing and processing large amounts of measurement data? Which technologies were found to be problematic? Are there new promising technologies that may be used in the future? What are the specific requirements for dealing with large-scale measurement data and how do they relate to or differ from other big data applications?
3. Which data analysis techniques have been found to be useful? Which data analysis techniques have been found to be problematic? Are there any novel promising techniques that need further research and development?

Although at the seminar the participants chose to organize the discussions on more general topics than these specific questions, during the discussions most of these questions were addressed to one degree or another.

## References

- 1 Flamingo – Management of the Future Internet. <http://www.fp7-flamingo.eu>. [Online; accessed 18-January-2016].
- 2 Leone – From Global Measurements to Local Management. <http://www.leone-project.eu>. [Online; accessed 18-January-2016].
- 3 RIPE Atlas: A Global Internet Measurement Network. *Internet Protocol Journal*, September 2015. URL: <http://ipj.dreamhosters.com/wp-content/uploads/2015/10/ipj18.3.pdf>.
- 4 V. Bajpai and J. Schönwälder. A Survey on Internet Performance Measurement Platforms and Related Standardization Efforts. *Communications Surveys Tutorials, IEEE*, 17(3):1313–1341, thirdquarter 2015. doi:10.1109/COMST.2015.2418435.
- 5 Vaibhav Bajpai, Steffie Jacob Eravuchira, and Jürgen Schönwälder. Lessons Learned From Using the RIPE Atlas Platform for Measurement Research. *SIGCOMM Comput. Commun. Rev.*, 45(3):35–42, July 2015. doi:10.1145/2805789.2805796.

- 6 kc claffy. The 7th Workshop on Active Internet Measurements (AIMS7) Report. *SIGCOMM Comput. Commun. Rev.*, 46(1):50–57, January 2016. doi:10.1145/2875951.2875960.
- 7 Philip Eardley, Marco Mellia, Jörg Ott, Jürgen Schönwälder, and Henning Schulzrinne. Global Measurement Framework (Dagstuhl Seminar 13472). *Dagstuhl Reports*, 2014. doi:  
<http://dx.doi.org/10.4230/DagRep.3.11.144>.
- 8 S. Sonntag, J. Manner, and L. Schulte. Netradar – Measuring the wireless world. WiOpt’13, May 2013.
- 9 Srikanth Sundaresan, Sam Burnett, Nick Feamster, and Walter De Donato. BISmark: A Testbed for Deploying Measurements and Applications in Broadband Access Networks. USENIX ATC’14, pages 383–394, 2014. URL: <http://dl.acm.org/citation.cfm?id=2643634.2643673>.
- 10 B. Trammell, P. Casas, D. Rossi, A. Bar, Z. Houidi, I. Leontiadis, T. Szemethy, and M. Mellia. mPlane: an Intelligent Measurement Plane for the Internet. *IEEE Communications Magazine*, 52(5):148–156, May 2014. doi:10.1109/MCOM.2014.6815906.

## 2 Table of Contents

### Executive summary

*Vaibhav Bajpai, Arthur W. Berger, Philip Eardley, Jörg Ott, Jürgen Schönwälder* . 15

### Invited Presentations

Experiences from Measuring Networks (Henning Schulzrinne) . . . . . 19

Empirical Network Science (Daniel Karrenberg) . . . . . 20

Global Measurements at Akamai (Arthur Berger) . . . . . 20

### Parallel Group Work

Measurement Platforms Integration . . . . . 21

Doing it Wrong . . . . . 22

Ethics . . . . . 23

Reproducibility and Data Quality . . . . . 24

Storage, Processing and Archival . . . . . 25

Future Measurement Challenges . . . . . 25

### Lightning Talks

HostView: Measuring Internet quality of experience on end-hosts (Renata Teixeira) 28

A Path Transparency Observatory (Brian Trammell) . . . . . 28

Virtual Measurement Accuracy (Al Morton) . . . . . 29

From Local to Global Measurements (Georg Carle) . . . . . 29

From Packet Counts to QoE (Markus Fiedler) . . . . . 29

CheesePi – Swedish home network monitoring (Ian Marsh) . . . . . 30

WebRTC Service Quality in the Wild (Varun Singh) . . . . . 30

Haystack – Mobile traffic monitoring in user space (Srikanth Sundaresan) . . . . . 31

Schengen Routing: A Compliance Analysis (Burkhard Stiller) . . . . . 31

**Conclusions and Next Steps** . . . . . 31

**Participants** . . . . . 33

### 3 Invited Presentations

The invited presentations were intended as a basis for triggering discussions and identifying areas for group work.

#### 3.1 Experiences from Measuring Networks (Henning Schulzrinne)

Henning Schulzrinne (Columbia University / FCC) began by sharing experiences gained through five iterations of the Federal Communications Commission (FCC) Measuring Broadband America (MBA) program [1], consisting of around 5,500 measurement hosts. The project is unique in that it is a collaboration between a regulator, a contractor (SamKnows [2]) developing and managing the infrastructure, about a dozen consumer ISPs and their trade associations, backbone ISPs, two third-party measurement facilities (M-Lab [3] and Level3) and university collaborators. Establishing a code of conduct and setting up a (light-weight) collaborative structure early on has helped work through conflicts and deal with data quality challenges. Since these measurements are used by competing providers, e.g., in TV commercials, the stakes are perceived to be higher than just scientific discovery. The project emphasizes long-term comparability of measurements, open data and reproducibility. For example, all scripts and spreadsheets used to produce the annual report are made available.

He described how the measurement report has changed, increasingly emphasizing variability in performance, across time and the user population, not just averages. One of the more contentious issues has been dealing with unexpected soft and subtle failures of measurement infrastructure, e.g., memory leaks and Ethernet port speed issues, as well as what to consider outliers. For example, a time period was excluded from the measurement month used for reporting since it coincided with the download traffic of iOS 8.0. Users may also delay upgrading their cable modem, causing performance to drop below the offered rate. In the long term, the current model of deploying hardware to end users does not scale well. It is hoped that building in-measurement functionality, e.g., through the IETF LMAP effort [2], rather than bolting it on later, will make measurement cheaper and more fine-grained.

He also emphasized that network diagnostics and network measurements can be highly complementary. For example, the ability to diagnose network problems may motivate end users to install network measurement devices and software. His recent research at Columbia University on measuring performance of YouTube streaming videos [4] finds a close correlation between QoE impairments and the abandonment of YouTube videos.

#### References

- 1 Measuring Broadband America – Federal Communications Commission. <https://www.fcc.gov/general/measuring-broadband-america>. [Online; accessed 18-January-2016].
- 2 V. Bajpai and J. Schönwälder. A Survey on Internet Performance Measurement Platforms and Related Standardization Efforts. *Communications Surveys Tutorials, IEEE*, 17(3):1313–1341, thirdquarter 2015. doi:10.1109/COMST.2015.2418435.
- 3 Constantine Dovrolis, Krishna Gummadi, Aleksandar Kuzmanovic, and Sascha D. Meinrath. Measurement Lab: Overview and an Invitation to the Research Community. *SIGCOMM Comput. Commun. Rev.*, 40(3):53–56, June 2010. doi:10.1145/1823844.1823853.
- 4 Hyunwoo Nam, Kyung-Hwa Kim, and Henning Schulzrinne. QoE Matters More Than QoS: Why People Stop Watching Cat Videos. In *INFOCOM, 2016 (to appear)*, April 2016.

### 3.2 Empirical Network Science (Daniel Karrenberg)

Daniel Karrenberg (RIPE NCC) shared experiences with doing empirical network science and derived some principles for good working practices from those experiences. He began by underlining the importance of reproducibility as a necessary condition for producing scientific work. In order to enable reproducibility, it requires one to archive everything during a scientific process. During an experiment, observations must be collected as close to the wire as possible. To avoid any mutation, the raw data derived from these observations must be archived with as little processing as possible. Virtual machines to build any software necessary should be encouraged. A good archive should also include documentation of the experiment such as immutable observations, metadata, experimental conditions, lab notes, calibration data, processed data, changelogs, comments, and analysis / publication backends. For instance, the experimental conditions must not only describe the state of the experiment but also document firmware/software versions to allow proper calibration. Since metadata (such as IP geolocations, IP reverse DNS records, IP prefix to origin AS mappings) is dynamic and volatile, it must also be archived in 'near observation' time. He encouraged the community to invest in storage not only for long term archival, but also to maximise headroom for future measurement results. He illustrated how around 2.1 PB of Hadoop Distributed File System (HDFS) [1] storage is currently (as of January 2016) allocated (with around 400 TB in use) for archiving measurement data produced by the RIPE Atlas project. He reasoned that a well organised and maintained archive not only makes analysis and publication easy, but also enables reuse of observations in the long run. Daniel related that data from some experiments in the RIPE Atlas public archive have indeed been re-used for different purposes. Furthermore providing basic controls to end-users enables unforeseen use of the measurement infrastructure as can be seen from creative usages of the RIPE Atlas measurement platform today.

#### References

- 1 Tom White. *Hadoop – The Definitive Guide: Storage and Analysis at Internet Scale (4. ed., revised & updated)*. O'Reilly, 2015. URL: <http://www.oreilly.de/catalog/9781491901632/index.html>.

### 3.3 Global Measurements at Akamai (Arthur Berger)

Based on experiences with the global measurement platform of Akamai Technologies, and relevant to the suggested topics for the seminar, Arthur Berger (Akamai Technologies) discussed three example performance metrics: (a) active measurements of latency and loss, which is used in Request Routing [1] to pick the best datacenter from which to serve a given client, (b) active measurements of latency and loss between Akamai servers, which is used to determine the via nodes in the Akamai routing overlay network [2] and (c) passive measurements of video downloads to clients, which shows the likelihood of abandonment by the end-user as a function of start-up time [3], indexed by the subject category of the video, such as sports, news or religion. He then discussed Akamai's processing of passive measurements recorded in log lines which consists of 1.2 PB data generated per day. He demonstrated where to find publicly available measurement results on Akamai's website [4]. Lastly, he suggested that an area that deserves more attention by the Internet measurement community is security and gave examples of some security measurements collected by Akamai.

## References

- 1 Fangfei Chen, Ramesh K. Sitaraman, and Marcelo Torres. End-User Mapping: Next Generation Request Routing for Content Delivery. SIGCOMM'15. ACM, 2015. doi: 10.1145/2785956.2787500.
- 2 Ramesh K. Sitaraman, Mangesh Kasbekar, Woody Lichtenstein, and Manish Jain. *Overlay Networks: An Akamai Perspective*, pages 305–328. John Wiley & Sons, Inc., 2014. doi: 10.1002/9781118909690.ch16.
- 3 S. Shunmuga Krishnan and Ramesh K. Sitaraman. Video Stream Quality Impacts Viewer Behavior: Inferring Causality Using Quasi-experimental Designs. IMC'12, 2012. doi: 10.1145/2398776.2398799.
- 4 State of the Internet – Akamai. <https://www.stateoftheinternet.com>. [Online; accessed 18-January-2016].

## 4 Parallel Group Work

The afternoon sessions were used to discuss certain topics in more depth in smaller groups. This section summarises the discussions of each group.

### 4.1 Measurement Platforms Integration

Lately there has been a rise in new and upcoming active measurement platforms [4] on a more or less equivalent underlying substrate (Linux on small cheap boxes). It is unclear whether one can (or should) integrate the common parts of these platforms. There are legitimate reasons to have diversity. For instance, each platform is designed with a distinct goal and provides separate coverage of sources to measure from. However, there is a lot of hard but repetitive work to create a new platform and keep it working. It is also unclear whether all platforms measure the basic measurement primitives in the same way, or whether there are some differences. If we knew they all worked in the same way, then we would be able to compare their results and potentially perform combined studies for a more comprehensive study. Furthermore, if the (common) test code was publicly available, then future developers would not need to expend efforts towards developing yet another version of the same measurement primitive. As such, the idea is to start by building a common codebase / measurement OS distribution for building an integrated measurement platform. The ingredients of this common codebase can include basic measurement utilities and package management tools.

A large number of use cases are covered by a few measurement primitives: (a) loss and latency using `ping`, (b) data-plane topology using `traceroute`, and (c) HTTP GET for applications. It is assumed that these primitives work the same everywhere for comparability reasons. However, there is a need for cross-calibration studies to confirm this premise. A meta-API that glues APIs from multiple measurement platforms together would allow studies to include vantage points from multiple platforms. A design of a Domain Specific Language (DSL) over these primitives implemented by the common substrate would further reduce the barrier to entry. A literature survey is needed to determine how much work such a common platform can save or how it would allow a measurement study to scale up. Measurements also come with a prerequisite for data storage and archival. A volunteer cloud for storage of measurement results with data replication could help spread care and feeding labour and ensure cross-institutional continuity of measurement results.

There are a number of challenges when integrating multiple measurement platforms. For one, reconciling design philosophies (simple vs. complex) is tricky. Seattle [6] is such an integrated measurement platform, albeit designed with a different goal to foster educational cloud computing but with a similar idea. However, the platform turns out to be an overkill for some simple use case scenarios. As such, a requirements description on necessary items for a minimal viable prototype is needed. It is also unclear how to form a community around this goal. It certainly helps to make participants feel good about doing something beneficial for the Internet but having venues to disseminate experience helps bring more people on board.

The management of the system is a first step towards integration. A vanilla OS distribution with stripped down packages with additional cross-compiled packages provided as an overlay similar to the BISmark platform [19] would be ideal. The goal should also be to allow the measurement suite to run inside a virtual machine. Virtual environments help keep dependency issues to a minimum. Given probes are remotely managed, another challenge is to avoid an update that renders them permanently unusable. BISmark uses a manual firmware update process with a possibility to fall-back to a trusted image to help mitigate this risk. Access control is another issue. It is unclear how much control the probe host must receive for hosting the probe.

The second step is to identify the measurement primitives that must be supported. Some candidate primitives may include: `dig`, `ping`, `curl`, `iperf` and a constant-bitrate packet generation tool. Including multiple variations of one primitive (such as `traceroute`) that are designed with slightly different goals (such as `scamper` [14] and `tracebox` [8]) adds value. The possibility of hosting multiple versions of each primitives must be supported. The primitives themselves must also support a common machine-readable output format. The ability of the primitives to write to a database (such as `sqlite`) increases the possibility of reuse since the results can simply be queried. It is a challenge to provide an exhaustive list of primitives that satisfies all measurement studies. As such requirements gathering and a survey to scope the problem is needed. For instance, `tcpdump` may be useful, but it has privacy implications and shipping data produced by this primitive may be a sensitive issue. A survey to identify incremental benefits of each measurement primitive is needed. It is also unclear if exceptions must be made for certain primitives to run in root privilege mode.

Furthermore, a number of future challenges were identified. For instance, a bootstrapping mechanism to get clients registered, an authentication mechanism to identify the clients, a server mechanism to be a destination for primitives, a communication channel to describe this client and server communication, an API to interface with measurement data, encryption and handling of key distribution are few identified areas that require work.

## 4.2 Doing it Wrong

A basic understanding of the strengths and weaknesses of a measurement method is useful since some weaknesses may inhibit interpretation of certain data and may lead to wrong conclusions. As such it is better to enumerate all the ways to collect the data one needs to answer a research question and then document their pros and cons. Continuous validation [13] is also important to produce data reliably, in particular if there is a dependency on 3rd party components (that may change in unanticipated ways). The key is to ask the question: Why do I have outliers or unexpected results? In order to be able to answer this question it is essential to have a world model [18] and some expectations of the data. At the same time one



also needs to be prepared for the world model to be wrong. Unexpected data requires careful analysis in order to determine whether there is a measurement error, a data analysis error, or a world model error. It is vital to be extensive in the description of the metadata [17] and the documentation of the experiment. As such, it is best to try to gather as much metadata (or context) as possible of the data, but at the same time also being honest about the limits of the data. One also needs to think in multiple timescales since time itself is a complicated thing to get right. Dealing with time can be notoriously hard due to accuracy and precision issues, clock drift issues, synchronization issues, issues caused by non-monotonic clocks and issues with time interpretation (such as ordering, timezone knowledge). Keeping raw data is important and so is the ability to reproduce the analysis.

### 4.3 Ethics

There can be a tension between scientific principles (measurements and meta-data should be public) and consistency with ethical principles. For instance, we have recently witnessed controversial papers [5, 10] that, although published, raised ethical concerns within the program committee.

There has already been activity by the community on ethical practice. For instance, a dedicated SIGCOMM workshop on Ethics in Networked Systems Research [1] was recently organized in 2015. Moreover, the call for papers for Internet Measurement Conference (IMC) encourages authors when appropriate to include a subsection describing ethical considerations and provides appropriate links for further information on ethical principles [11] and guidance [2] on ethical data sharing. As a followup to the Dagstuhl seminar on Ethics in Data Sharing [7, 9], SURFnet is preparing a document [21] on Data Sharing policy. The final policy will most likely come into effect in the first quarter of 2016.

The issue of what can be considered ethical is often a grey area since opinions can dramatically vary by different parties. For example, a study [3] that analyses causes of collateral damage of censorship by identifying DNS injection activities of the Great Firewall of China could potentially be viewed as unethical by the government of the People's Republic of China. Several intriguing questions from the ethical standpoint deserve discussion. For instance, in a measurement study of cyber crime, is it appropriate to buy products from criminals? and is it appropriate to crawl a website to obtain all of the information even when the site explicitly states that one should not do this?

Ethical issues pertain to more than just privacy infringement. For instance, disrupting the service of an end-user and possibly even endangering an end-user without the user's consent. There is a fine line between legal and ethical issues. Moreover, ethical issues encompass the entire measurement chain starting from the design of an experiment, conducting measurements, data storage, data processing, and data sharing. The security research community is increasingly sensitive to this issue. For instance, some IT departments avoid collecting data just so that they have no data if asked by a law-enforcement agency. By nature, research tends to push the boundaries, however risk analysis can be hard. If it is known in advance how the data will be used, collect just what is needed.

The Internet measurement community needs to publish further guidelines on ethical practice. A key target audience is researchers that are not aware of the issues, but would want to do the right thing. For the Internet measurement community, continued discussion to gain more clarity in the aforementioned grey areas is needed. Regardless of whether there is consensus in the community as a whole, a conference program committee should have the

discretion to reject a paper on ethical grounds. The authors of the rejected paper could be asked for permission to make known the aspect of the work that was considered unethical, so as to provide guidance to the wider community. Furthermore, since a paper rejection occurs after the unethical practice has already occurred, the goal must be to avoid the unethical practice to happen in the first place. To address this part, an interesting question is how to have curricula embrace ethics educations. An ethics background is not just only needed for measurement studies, but in general for people working in computer science, both in academia and industry.

#### 4.4 Reproducibility and Data Quality

Repeatability and reproducibility are often misinterpreted in practice. Repeatability is the notion of re-running the same experiment with a change in time. Reproducibility on the other hand is being able to derive the same conclusions with a change in both space and time. Reproducibility imbibes the flexibility of using different measurement methods to arrive at the same conclusion. In order to foster reproducibility, a number of aspects need to be documented: (a) measurement method, (b) metric, (c) vantage points and (d) implementation. This requires a characterisation plan of statistical tests to imply significance of data analysis.

There are a number of difficulties in reproducing an experiment. For one, statistical analysis is hard. There is a danger of confirmation bias with a tendency to abandon experiments if results are boring. It is often not possible to measure the metric directly (or only as a one-off calibration) since generally there is a lack of stable ground truth. Moreover, it is difficult to publish a study that reproduces an experiment. Worse, documenting the limitations of an experiment is often (wrongly) seen as a weakness. Particularly, sharing datasets of an experimental study with others is hard. For one, legal rules vary by jurisdiction but more so one wants to have a first mover advantage with the associated data collection activity.

There is a need to add rigour in statistical analysis to enable reproducibility. This starts with hypothesis testing and concrete research questions. Factor analysis during the experimental design to cover the design space of variables is often ignored. Outliers that fail hypothesis need treatment. As such, a prospective journal that invites reproducibility would help reduce probability of early abandonment of experiments that confirm previous results. Calibration and quality checks must be encouraged. Conferences can be encouraged to dedicate special sessions devoted to papers that reproduce results. Researchers on the other hand must be encouraged to write a technical report that describes the dataset used in a publication. Such a report must document the measurement method, dataset fields, limitations and scope of the dataset. There are QoE standards that provide test conditions and advice on where the measurement method is applicable. In cases where raw dataset cannot be shared for some reason, researchers must still be encouraged to share the dataset in at least some restricted form by either removing some data columns, obfuscating some fields, or by allowing limited access to the dataset using SQL queries.

## 4.5 Storage, Processing and Archival

There is currently lack of a best current practice guide on how to store, process and archive measurement data. It seems that academics generally tend to rely on a Network-attached Storage (NAS) coupled with a few highly performant data crunching machines for data analysis operations.

One clear recommendation is to transition away from NAS because they cannot provide local computing power. Apache Hadoop [23] is a better alternative since it provides a tight coupling of storage and compute power, scales gradually over time and in the process turns out to be cost effective. Cloudera Distribution Including Apache Hadoop (CDH) packages [23] provide a simple head start into the Hadoop ecosystem. They can be used to deploy the Hadoop cluster and packages provide tools to make management easy. A transition to Hadoop can be done in multiple iterations. A first step is to get it functional by storing already existing measurement data in CSV or JSON format. However, in the long run a good serialization format such as Apache Avro [23] (for row-oriented datasets) or Apache Parquet [23] (for columnar storage) can help future proof storage in a structured format. Naturally, it is better to make a choice at the very outset of data collection. HBase [23], a non-relational database that can run on top of HDFS, can be used for high performance analysis for specialised applications, although it tends to have a steeper learning curve especially for SQL users. Cloudera Impala adds an SQL engine on top of HDFS. GraphQL can be used to decouple presentation from querying on the data. Message queues such as RabbitMQ can be used for stream-based processing requirements. A recently developed large-scale active measurement platform for DNS [22] uses Hadoop for storage and analysis of data. Experiences with this platform show that there is potential for using Hadoop for Internet measurements.

Certainly the Hadoop tool chain is not the answer to all problems. It must be viewed as HDFS for storage with optional powerful processing on top. Although at times, a powerful machine with lots of computing cores can also serve the same task at hand. As such, at the end the size of the data matters. Hadoop distributes I/O and processing and thus it can crunch large volumes of data in short time. A single fast machine has I/O limits and may have CPU / memory limits that are difficult to scale (but it is often the I/O limit that is difficult or expensive to change).

## 4.6 Future Measurement Challenges

Measurement methods will evolve beyond traditional active and passive techniques. Al Morton in [15] describes hybrid measurement methods which are subset of both active and passive methods. For instance, Type I hybrid measurements employ methods that augment or modify the stream of interest, while Type II hybrid measurements employ methods that utilize two or more streams of interest with some degree of mutual coordination to collect multiple metrics.

A number of Internet measurement tools are designed with inherent assumptions (about layer-2 networks) [20] that are not true for underlying wireless links. In particular WiFi home networks and cellular networks are impacted by bitrates, retransmission rates, and signal strengths as wireless channel conditions change. As such, we need to design measurement approaches and tools that are also suitable for measuring wireless links.

There are also challenges with metrics that measure available bandwidth. In the view of mostly elastic traffic, partly in combination with wireless links, it is not clear whether a

convincing solution can be expected. Moreover, with existing tools, probing for capacity does not work well with tools that assume that the link is work-conserving.

For many web-based applications, end-to-end traffic is split [16] into a transport session from end system to the front-end servers, and another transport session to the backend infrastructure. In the transport session to the front-end servers, many short-term TCP flows may be observed (in contrast to long-lived TCP flows in the transport session to the backend infrastructure). In such a scenario, protocols used to establish the transport session to the front-end servers can be changed quickly. For instance, Quick UDP Internet Connections (QUIC) [12], which is increasingly used to establish a transport session to the front-end servers, may behave more aggressively than TCP.

There is an increasing demand for low-latency communication. Many technological advances reduce latency significantly. For instance, compared with 4G, 5G claims that it will reduce latency by a factor of 100. However, it is unclear how one can measure latency in these new environments with the required level of accuracy. The security aspects of Internet of Things (IoT) devices are becoming critical. It is unclear whether there is a need for specialized measurement tools and methods in this space. An analysis of measurement challenges with respect to IoT security is needed.

A large number of network functions are being virtualised today. It remains unclear how to measure in such virtualised scenarios. Additional measurement objectives and metrics need to be identified particularly due to the resource sharing effects of such virtual network functions.

## References

- 1 Workshop on Ethics in Networked Systems Research – ACM SIGCOMM 2015. <http://conferences.sigcomm.org/sigcomm/2015/netethics.php>. [Online; accessed 18-January-2016].
- 2 Mark Allman and Vern Paxson. Issues and Etiquette Concerning Use of Shared Measurement Data. IMC’07, pages 135–140. ACM, 2007. doi:10.1145/1298306.1298327.
- 3 Anonymous. The Collateral Damage of Internet Censorship by DNS Injection. *SIGCOMM Comput. Commun. Rev.*, 42(3):21–27, June 2012. doi:10.1145/2317307.2317311.
- 4 V. Bajpai and J. Schönwälder. A Survey on Internet Performance Measurement Platforms and Related Standardization Efforts. *Communications Surveys Tutorials, IEEE*, 17(3):1313–1341, thirdquarter 2015. doi:10.1109/COMST.2015.2418435.
- 5 Sam Burnett and Nick Feamster. Encore: Lightweight Measurement of Web Censorship with Cross-Origin Requests. SIGCOMM’15. ACM, 2015. doi:10.1145/2785956.2787485.
- 6 Justin Cappos, Ivan Beschastnikh, Arvind Krishnamurthy, and Tom Anderson. Seattle: A Platform for Educational Cloud Computing. SIGCSE’09. ACM, 2009. doi:10.1145/1508865.1508905.
- 7 Julie E. Cohen, Sven Dietrich, Aiko Pras, Lenore D. Zuck, and Hildebrand Mireille. Ethics in Data Sharing (Dagstuhl Seminar 14052). *Dagstuhl Reports*, 4(1):170–183, 2014. doi: <http://dx.doi.org/10.4230/DagRep.4.1.170>.
- 8 Gregory Detal, Benjamin Hesmans, Olivier Bonaventure, Yves Vanaubel, and Benoit Donnet. Revealing Middlebox Interference with Tracebox. IMC’13. ACM, 2013. doi:10.1145/2504730.2504757.
- 9 Sven Dietrich, Jeroen van der Ham, Aiko Pras, Roland van Rijswijk-Deij, Darren Shou, Anna Sperotto, Aimee van Wynsberghe, and Lenore Zuck. Ethics in Data Sharing: developing a model for best practice. CREDs II, 2014.
- 10 Marcel Dischinger, Andreas Haeberlen, Krishna P. Gummadi, and Stefan Saroiu. Characterizing Residential Broadband Networks. IMC’07. ACM, 2007. doi:10.1145/1298306.1298313.

- 11 D. Dittrich and E. Kenneally. The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. Technical report, U.S. Department of Homeland Security, 2012.
- 12 Ryan Hamilton, Jana Iyengar, Ian Swett, and Alyssa Wilk. QUIC: A UDP-Based Secure and Reliable Transport for HTTP/2. Internet-Draft draft-tsvwg-quic-protocol-02, Internet Engineering Task Force, January 2016. URL: <http://tools.ietf.org/html/draft-tsvwg-quic-protocol-02>.
- 13 Balachander Krishnamurthy, Walter Willinger, Phillipa Gill, and Martin Arlitt. A Socratic method for validation of measurement-based networking research. *Computer Communications*, 34(1):43–53, 2011. doi:<http://dx.doi.org/10.1016/j.comcom.2010.09.014>.
- 14 Matthew Luckie. Scamper: A Scalable and Extensible Packet Prober for Active Measurement of the Internet. IMC'10. ACM, 2010. doi:10.1145/1879141.1879171.
- 15 Al Morton. Active and Passive Metrics and Methods (and everything in-between, or Hybrid). Internet-Draft draft-ietf-ippm-active-passive-06, Internet Engineering Task Force, January 2016. URL: <http://tools.ietf.org/html/draft-ietf-ippm-active-passive-06>.
- 16 Abhinav Pathak, Y. Angela Wang, Cheng Huang, Albert Greenberg, Y. Charlie Hu, Randy Kern, Jin Li, and Keith W. Ross. Measuring and Evaluating TCP Splitting for Cloud Services. PAM'10, pages 41–50. Springer-Verlag, 2010. URL: <http://dl.acm.org/citation.cfm?id=1889324.1889329>.
- 17 Vern Paxson. Strategies for Sound Internet Measurement. IMC'04, pages 263–271. ACM, 2004. doi:10.1145/1028788.1028824.
- 18 M. Roughan, W. Willinger, O. Maennel, D. Perouli, and R. Bush. 10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems. *JSAC*, 29(9):1810–1821, October 2011. doi:10.1109/JSAC.2011.111006.
- 19 Srikanth Sundaresan, Sam Burnett, Nick Feamster, and Walter De Donato. BISmark: A Testbed for Deploying Measurements and Applications in Broadband Access Networks. USENIX ATC'14, pages 383–394, 2014. URL: <http://dl.acm.org/citation.cfm?id=2643634.2643673>.
- 20 Srikanth Sundaresan, Nick Feamster, and Renata Teixeira. Measuring the Performance of User Traffic in Home Wireless Networks. PAM'15, pages 305–317, 2015. doi:10.1007/978-3-319-15509-8\_23.
- 21 Roland van Rijswijk-Deij. Ethics in Data Sharing: a best practice for NRENs. TNC'15. GÉANT, 2015. URL: <https://tnc15.terena.org/getfile/1869>.
- 22 Roland van Rijswijk-Deij, Mattijs Jonker, Anna Sperotto, and Aiko Pras. The Internet of Names: A DNS Big Dataset. SIGCOMM'15, pages 91–92. ACM, 2015. doi:10.1145/2785956.2789996.
- 23 Tom White. *Hadoop – The Definitive Guide: Storage and Analysis at Internet Scale (4. ed., revised & updated)*. O'Reilly, 2015. URL: <http://www.oreilly.de/catalog/9781491901632/index.html>.

## 5 Lightning Talks

Participants were also encouraged to volunteer for a lightning talk to provide a perspective into their recent measurement research work.

### 5.1 HostView: Measuring Internet quality of experience on end-hosts (Renata Teixeira)

There is interest in automated performance diagnosis on user laptops or desktops. One interesting aspect that has received little attention is the user perspective on performance. To conduct research on both end-host performance diagnosis and user perception of network and application performance, Renata Teixeira (INRIA) presented an end-host data collection tool, called HostView. HostView [1] not only collects network, application and machine level data, but also gathers feedback directly from users. User feedback is obtained via two mechanisms: a system-triggered questionnaire and a user-triggered feedback form. In her talk, she described experiences with the first deployment of HostView. Using data from 40 users, she articulated challenges in this line of research, and reported initial findings in correlating user data to system-level data. She then described more recent efforts in conducting an in-depth study with 12 users in France to guide the design of the next version of HostView and of methods to infer user context and activities.

#### References

- 1 Ahlem Reggani, Fabian Schneider, and Renata Teixeira. An End-Host View on Local Traffic at Home and Work. PAM'12, pages 21–31, 2012. doi:10.1007/978-3-642-28537-0\_3.

### 5.2 A Path Transparency Observatory (Brian Trammell)

The growing deployment of middle boxes in the Internet has reduced the degree to which the Internet is still an end-to-end network in accordance with its original design. This lack of end-to-endness leads to ossification of the transport layer [1]: new protocols are difficult or impossible to deploy as they must be designed around middle boxes, either those which have been observed, or conjectured to exist. It is necessary to guide protocol engineering for transport protocol innovation on a basis of observations of the Internet as it is, but these observations are hard to come by. Brian Trammell (ETH Zürich) proposed a Path Transparency Observatory, which can take observations of path transparency (the likelihood a packet stream that arrives at the end of the path is the one that was sent, with certain properties) and impairment (something that keeps a path from being transparent for a certain kind of traffic) from multiple sources, with multiple resolutions of condition definition and information about the endpoints and path involved. An observatory collects single observations of a path and a condition on that path at some point in time, with references to the code that created the observations so they can be repeated, and a set of equivalence functions so that equivalent conditions and paths can be compared. He explained that this work is ongoing, and a public observatory will become available within the scope of the Measurement and Architecture for a Middleboxed Internet (MAMI) project [2] over next two and a half years.

#### References

- 1 Michio Honda, Yoshifumi Nishida, Costin Raiciu, Adam Greenhalgh, Mark Handley, and Hideyuki Tokuda. Is It Still Possible to Extend TCP? IMC'11. ACM, 2011. doi:10.1145/2068816.2068834.
- 2 MAMI – Measurement and Architecture for a Middleboxed Internet. <https://mami-project.eu>. [Online; accessed 18-January-2016].

### 5.3 Virtual Measurement Accuracy (Al Morton)

The movement towards Network Function Virtualization (NFV) means that measurement system virtualisation will take place for active, passive, and hybrid methods of measurement. This evolution will allow on-demand deployment of measurement systems in general purpose servers. The designs must be cost-effective, but there is tension between cost of physical resources and accuracy. Al Morton (AT&T) presented this trade-off and associated challenges.

### 5.4 From Local to Global Measurements (Georg Carle)

Georg Carle (TU München) provided a summary of measurement based research work conducted within his group. He explained that understanding Internet phenomena requires both local and global measurements. Local measurements such as on the MEMPHIS test bed allows for reproducible experiments. As part of this project, the MoonGen Traffic generator [1] is an example that allows for high precision by directly accessing hardware features such as precise time stamping from the application space, while bypassing the operating system. Furthermore, Georg reasoned how Software-defined Network (SDN) mechanisms can be used for performing very high-speed flow monitoring using Commercial Off-the-shelf (COTS) components and adaptive load-balancing. One objective of security-related global measurements is to identify prefix hijacking. An innovative approach [2] to identify benign anomalies is to use information with business relations and ownership information from a publicly accessible Internet Routing Registry (IRR) to combine it with collected TLS certificates, and using these certificates as fixed points to be checked in time intervals in which routing anomalies are observed. For performing measurements with wireless links, he presented the MearDroid Android app, which allows to perform wireless measurements from many vantage points.

#### References

- 1 Paul Emmerich, Sebastian Gallenmüller, Daniel Raumer, Florian Wohlfart, and Georg Carle. MoonGen: A Scriptable High-Speed Packet Generator. IMC'15. ACM, 2015. doi: 10.1145/2815675.2815692.
- 2 Johann Schlamp, Ralph Holz, Oliver Gasser, Andreas Korsten, Quentin Jacquemart, Georg Carle, and Ernst W. Biersack. Investigating the Nature of Routing Anomalies: Closing in on Subprefix Hijacking Attacks. TMA'15. doi:10.1007/978-3-319-17172-2\_12.

### 5.5 From Packet Counts to QoE (Markus Fiedler)

Markus Fiedler (BTH) discussed challenges in measuring QoE. He stressed that the main challenge for the interpretation of network measurements in light of QoE is that user perception happens far up the network stack, far away from where Quality of Service (QoS) problems (such as latency and packet loss) arise and where monitoring takes place. QoS may be transformed significantly throughout the stack. The recently proposed QoE Hourglass Model [1] is one way to formalize such transformations, capturing impacts of transport protocols, display devices and other factors. Using an example from a project with a major European telecommunications provider, he proposed a method that enables the exploitation of information from packet counters for QoE assessment. It starts with the definition of the user-perceived problem to be attacked, followed by the determination of parameters that



reflect those problems far down in the network stack, and of the critical timescale for the user, and finally the use of appropriate comparative summary statistics.

#### References

- 1 Tahir Nawaz Minhas and Markus Fiedler. Quality of experience hourglass model. In *ComManTel 2013*, Jan 2013. doi:10.1109/ComManTel.2013.6482371.

### 5.6 CheesePi – Swedish home network monitoring (Ian Marsh)

Ian Marsh (SICS) described the architecture of a distributed measurement system, CheesePi. He utilized the IETF LMAP framework [1] terminology to describe this system. CheesePi uses Raspberry Pi hardware devices to allow always-on, simple and reliable monitoring of users' home Internet connections. By running CheesePi on a Raspberry Pi (termed a Measurement Agent (MA)) connected to their home network, a non-expert user can continuously monitor their connection quality. He argued that a common hardware platform for all MAs gives greater consistency between the collected measurements and it also simplifies the codebase. The result is a common software platform for measurement tasks that can host, execute and record arbitrary network behaviour. Ian explained the reason for deploying dedicated monitoring devices. The project is tailored towards capturing the network connectivity that devices are able to achieve. This can significantly depend on the last hop technology (e.g. Ethernet or WiFi), which would be missed by passive monitoring of user traffic at the home gateway. This work is performed in collaboration with the Swedish regulator Post and Telecom Authority (PTS), who are particularly concerned with expanding connection performance metrics from naive throughput measurements of a particular location and time to something more instructive. He argued that an easily comparable and widely understood metric (e.g., download/upload rates) does not necessarily indicate the QoE of a user.

#### References

- 1 P. Eardley, A. Morton, M. Bagnulo, T. Burbridge, P. Aitken, and A. Akhter. A Framework for Large-Scale Measurement of Broadband Performance (LMAP). RFC 7594 (Informational), September 2015. URL: <http://www.ietf.org/rfc/rfc7594.txt>.

### 5.7 WebRTC Service Quality in the Wild (Varun Singh)

Varun Singh (Aalto University) introduced callstats.io, a Web Real-Time Communication (WebRTC) analytics and diagnostics service. It measures service- and conference-level metrics for a WebRTC application service. At the service-level, annoyances (such as, how often do conferences fail, what are the reasons for failure and what is the typical network latency?) are measured. Varun described how callstats.io will share the aggregate quality metrics measured across tens of WebRTC services (big and small, local and global) with the measurement community at large.



## 5.8 Haystack – Mobile traffic monitoring in user space (Srikanth Sundaresan)

Despite our growing reliance on mobile phones for a wide range of daily tasks, we remain largely in the dark about the operation and performance of our devices, including how (or whether) they protect the information we entrust to them, and with whom they share it. The absence of easy, device-local access to the traffic of our mobile phones presents a fundamental impediment to improving this state of affairs. To develop detailed visibility, Srikanth Sundaresan (ICSI) presented Haystack [1], a system for unobtrusive and comprehensive monitoring of network communications on mobile phones, entirely from user-space. Haystack correlates disparate contextual information such as app identifiers and radio state with specific traffic flows destined to remote services, even if encrypted. Haystack facilitates user-friendly, large-scale deployment of mobile traffic measurements and services to illuminate mobile app performance, privacy and security. Srikanth described the design of Haystack and demonstrated its feasibility with an implementation that provides 26-55 Mbps throughput with less than 5% CPU overhead. He stressed that the system and results highlight the potential for client-side traffic analysis to help understand the mobile ecosystem at scale.

### References

- 1 Abbas Razaghpanah, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Christian Kreibich, Phillipa Gill, Mark Allman, and Vern Paxson. Haystack: In Situ Mobile Traffic Analysis in User Space. *CoRR*, abs/1510.01419, 2015. URL: <http://arxiv.org/abs/1510.01419>.

## 5.9 Schengen Routing: A Compliance Analysis (Burkhard Stiller)

Burkhard Stiller (UZH) described Schengen routing as a strategy to keep traffic originating from sources located in the Schengen area (an area comprising of 26 European countries that have abolished passport and any other type of border control at their common borders) and targeted to destinations located in the Schengen area within the Schengen area. He summarised results of a larger-scale measurement effort [1] performed to quantify Schengen routing compliance in parts of today's Internet. Based on a few thousand TCP, UDP, and ICMP `traceroute` measurements executed from RIPE Atlas probes located in over 1100 different Autonomous Systems (AS) in the Schengen area, it was observed that 34.5% to 39.7% of these routes are Schengen-compliant, while compliance levels vary from 0% to 80% among countries.

### References

- 1 Daniel Dönni, Guilherme Sperb Machado, Christos Tsiaras, and Burkhard Stiller. Schengen Routing: A Compliance Analysis. In *AIMS 2015, Ghent, Belgium, June 22-25, 2015.*, 2015. doi:10.1007/978-3-319-20034-7\_11.

## 6 Conclusions and Next Steps

Participants with a mix of senior and junior researchers hailing from both academia and industry encouraged fruitful dialogue. A number of future research agendas were recognized. Brian Trammell volunteered to initiate further discussion on the seminar mailing list towards measurement platform integration. An action item to create a code repository to hold basic primitives that can output results in a machine readable manner was created. Furthermore,

discussion on an Internet measurement cloud for not only storing measurement results but also facilitate its reliable distribution will begin. The organizing team also received valuable feedback. An interest to identify a specific problem to try to tackle it during a prospective future seminar was identified.

### **Acknowledgements**

The organisers would like to thank the participants for their contributions: Special thanks to Olivier Bonaventure, Henning Schulzrinne, Ian Robin Marsh, and Roland van Rijswijk-Deij for reviewing the manuscript. This work was funded by Flamingo, a Network of Excellence project (ICT-318488) supported by the European Commission under its Seventh Framework Programme.

## Participants

- Vaibhav Bajpai  
Jacobs University – Bremen, DE
- Arthur W. Berger  
Akamai Technologies –  
Cambridge, US
- Georg Carle  
TU München, DE
- Renata Cruz Teixeira  
INRIA – Le Chesnay, FR
- Philip Eardley  
BT Research – Ipswich, GB
- Markus Fiedler  
Blekinge Institute of Technology –  
Karlskrona, SE
- Phillipa Gill  
Stony Brook University, US
- Oliver Hohlfeld  
RWTH Aachen, DE
- Steffie Jacob Eravuchira  
SamKnows Ltd. – London, GB
- Daniel Karrenberg  
RIPE NCC – Amsterdam, NL
- Mirja Kühlewind  
ETH Zürich, CH
- Andri Lareida  
Universität Zürich, CH
- Jukka Manner  
Aalto University, FI
- Ian Robin Marsh  
Swedish Institute of Computer  
Science – Kista, SE
- Al Morton  
AT&T – Middletown, US
- Jörg Ott  
TU München, DE
- Colin Perkins  
University of Glasgow, GB
- Philipp Richter  
TU Berlin, DE
- Jair Santanna  
University of Twente, NL
- Jürgen Schönwälder  
Jacobs University – Bremen, DE
- Henning Schulzrinne  
Columbia Univ. – New York, US
- Varun Singh  
Callstats.IO, FI
- Burkhard Stiller  
Universität Zürich, CH
- Srikanth Sundaresan  
ICSI – Berkeley, US
- Brian Trammell  
ETH Zürich, CH
- Roland van Rijswijk-Deij  
University of Twente, NL

