

The 2nd Workshop on Active Internet Measurements (AIMS-2) Report

kc claffy
CAIDA
kc@caida.org

Emile Aben
RIPE NCC
emile.aben@ripe.net

Jordan Augé
UPMC Paris Universit s
jordan.auge@lip6.fr

Robert Beverly
Naval Postgraduate School
rbeverly@nps.edu

Fabian Bustamante
Northwestern University
fabianb@eecs.northwestern.edu

Benoit Donnet
Universit  catholique de
Louvain / FNRS
benoit.donnet@uclouvain.be

Timur Friedman
UPMC Paris Universit s
timur.friedman@upmc.fr

Marina Fomenkov
CAIDA
marina@caida.org

Peter Haga
Eotvos Lorand University
haga@elte.hu

Matthew Luckie
University of Waikato
mjl@wand.net.nz

Yuval Shavitt
Tel Aviv University
shavitt@eng.tau.ac.il

ABSTRACT

On February 8-10, 2010, CAIDA hosted the second Workshop on Active Internet Measurements (AIMS-2) as part of our series of Internet Statistics and Metrics Analysis (ISMA) workshops. The goals of this workshop were to further our understanding of the potential and limitations of active measurement research and infrastructure in the wide-area Internet, and to promote cooperative solutions and coordinated strategies to addressing future data needs of the network and security research communities. The three-day workshop included presentations, group discussion and analysis, and focused interaction between participating researchers, operators, and policymakers from all over the world. This report describes the motivation and findings of the workshop, and reviews progress on recommendations developed at the 1st Active Internet Measurements Workshop in 2009 [18]. Slides from the workshop presentations are available at [9].

1. MOTIVATION

The 1st AIMS Workshop in 2009 [8] identified relevant stakeholders who may support and/or oppose measurement, and explored how collaborative solutions might maximize the collective benefit of deployed infrastructure and gathered measurements, especially with respect to maintaining longitudinal coverage. The final report from that workshop [18] outlined open research problems identified by participants, and issued recommendations that could benefit both Internet science and communications policy. These recommendations represent a multi-year roadmap of the landscape with specific suggestions for paths to advance the quality, science, and utility of active Internet measurements. The AIMS workshop series provides a forum to track, evaluate, and build on previous achievements, refine our understanding of remaining problems and recognize new ones, and modify the course of progress as necessary.

The second Active Internet Measurements (AIMS-2) workshop convened on February 8-10, 2010, focused on continuing interaction among international researchers involved in active measurements of the Internet as well as exchange of insights among researchers, operators, and policymakers. The Internet research community develops novel infrastructures, as well as measurement and analysis methodologies, to inform our understanding of Internet structure, behavior, performance, and surrounding

communications policy, and to serve as input and validation for emerging Internet models. In the last decade, as interest in critical infrastructure protection has provided additional motivation for understanding fine-grained Internet topologies, we have also improved our understanding of technical, legal, ethical and policy issues that constrain the scope, availability, and usability of such infrastructure.

The workshop achieved its basic goals: participants presented their ideas, techniques and findings, and openly discussed how to navigate issues inhibiting progress, including how projects could leverage each other's infrastructures. Topics discussed included geolocation, infrastructure platforms, topology measurement and mapping techniques including for IPv6, alias resolution, end-to-end performance measurement, and data sharing needs and approaches. This report summarizes key findings and status updates regarding recommendations from last year's workshop.

2. KEY FINDINGS

The field of Internet mapping is making slow but steady progress.

Active measurements of the Internet have been applied to discover, model, and generate realistic macroscopic Internet topologies at various granularity levels: IP addresses, routers, autonomous systems (ASes). Improvements in measurement, inference, and validation methods, as well as the increasing availability of historical meta-data, have allowed quantitative as well as qualitative improvement in the best available Internet topology maps [15, 21, 17, 10]. More researchers are pursuing sophisticated descriptions of Internet topologies, annotated with performance and economic attributes: inter-AS relationships, bandwidth, RTTs, DNS information, IP-to-router and router-to-AS assignments, etc. These more richly annotated maps promise more realistic and comprehensive representations of available data, yielding more accurate representations and reflections of the global Internet, helping researchers to avoid some of the known pitfalls of infrastructure based measurements from limited vantage points [11].

Geolocation, i.e., determining geographical positions of Internet identifiers (IP addresses, domain names, AS numbers) is essential to many research investigations.

Governments, researchers, and commercial entities all seek improved capabilities to map Internet resources to geophysical locations. Despite the high level of demand, the field of geolocation

is still a sub-scientific discipline. While commercial IP geolocation tools exist [7], they tend to use proprietary methodologies, offer poor granularity, often disagree with each other in location results, and rarely include geolocation functionality for IPv6 addresses. There are no commonly accepted standard protocols, data formats, quality of service gauges, or comparison of existing tools and services. Some groups (W3C, IETF GeoPriv WG) are trying to standardize on interfaces to support location-aware Internet services, [31] [3], but progress is slow. Without ground truth data, measurement-based geolocation [23] can support limited validation of other geolocation estimation methods, as reported shortly after the workshop in a DIMES study [29]. Participants agreed that funding agencies should provide targeted support to stimulate more work in geolocation, both for wired and wireless applications, as well as evaluation and performance comparison of what geolocation tools and services exists today. CAIDA solicited input on a geolocation tool comparison study they are planning for this year as part of their topology mapping project [5].

There is still room for much improvement in interaction and technology transfer between three Internet research stakeholders: academic laboratories, commercial enterprises, and government institutions. In the absence of wide-scale cooperation from ISPs regarding sharing topology data, cross-correlation across independent results obtained by different measurement methods remains the most reasonable validation approach.

AIMS-2 participants reiterated a strong consensus articulated at previous workshops [22, 18] – that lack of ground truth data for validation of inferences continues to present the most serious bottleneck for rigorous development of Internet science. Informal contacts between researchers and individual providers can and do produce useful sets of ground truth data, but these datasets are limited in scope, not necessarily representative of larger topology maps, and not generally shareable, rendering it difficult to establish or verify the scientific integrity of published results. Furthermore, the validity of topology ground truth can quickly decrease with time, while many analyses need accurate ground truth data concurrent with executed measurements. Both drawbacks suggest the need for more scalable approaches to validation, but without more formal widespread cooperation from ISPs, improving validation will require closer collaboration across research and infrastructure projects, in particular exchange of tools, data, and measurement methodologies to allow cross-validation and reproducibility of results.

2.1 Infrastructure updates

One session of talks was devoted to operators of active measurement infrastructure projects to provide status updates regarding scope, functionality, and activities of supported platforms, with subsequent discussion of whether there are more effective ways to fund and cross-fertilize across infrastructure projects.

TopHat (<http://www.top-hat.info/>) is a measurement service to be made publicly available in September 2010 by UPMC Sorbonne Universit s, as part of the OneLab experimental facility. It is designed to support experiments on the PlanetLab testbed, from setup through completion. For example, TopHat gives PlanetLab users access to third-party measurement systems capable of higher precision timestamps, such as ETOMIC (<http://www.etomic.org/>), through a familiar PlanetLab-based API. TopHat has its own dedicated measurement infrastructure, TDMI, that runs in a PlanetLab slice and provides baseline measurements between PlanetLab nodes. TDMI conducts traceroutes using Paris Traceroute [30], to capture accurate paths through load-balancing routers. For measurements from PlanetLab vantage points, TopHat can also call upon a score of specialized ETOMIC boxes collocated with PlanetLab nodes. Since applications on PlanetLab often have a scope well beyond PlanetLab itself, vantage points outside of PlanetLab are also valuable to users. TopHat also interconnects with the DIMES platform (<http://www.netdimes.org/>), offering an interface to retrieve measurements from DIMES agents around the world.

iPlane [24, 25, 1] is a system developed by researchers at UW

for Internet-wide performance prediction. As of June 2010, iPlane runs across roughly 300 sites in the PlanetLab testbed and 200 public traceroute servers. iPlane continually issues measurements from all of these vantage points to more than 90% of prefixes at the Internet's edge and then uses measurements to make predictions of performance along unmeasured paths. First deployed in June 2006, iPlane has been used in both production systems and research projects at more than 50 institutions to date.

DIMES, an Internet topology measurement project based at Tel Aviv University, has increased its scope of measurements to include Planetlab nodes. As of June 2010, DIMES collects 3.5-4.5 million measurements a day from about 1500 agents.

Archipelago (Ark) [19] is CAIDA's active measurement platform providing a source of curated Internet topology and performance data as well as a virtual laboratory for independent experimentation. The laboratory includes a layer of "middleware" – specifically a coordination and communication facility to support macroscopic distributed asynchronous Internet measurements, shielding researchers from the complexities of managing measurement infrastructure. As of June 2010, Ark is composed of 45 nodes capable of flexible probing of IPv4 address space, 11 with IPv6 capability. Ark is now supporting IPv6 Spoofer measurements described at last year's AIMS workshop [4].

EdgeScope [2, 12, 13] is a project aimed to expose the network view of end systems located at the edge of the network. For this, it relies on the unique and valuable network perspective of peers in large-scale peer-to-peer systems, one to which today's researchers, network operators and users have limited or no access. Through extensions to a popular BitTorrent client, now in use by more than 1 million users in over 200 countries, EdgeScope collects tens of GB of network data daily, comprising data gathered from more than 4 million IP addresses located in over 8,700 ASes and 73,000 routable prefixes.

3. RESULTING COLLABORATIONS

Collaborations that have derived from the AIMS workshop have included improvements to Ark infrastructure to facilitate its use by other research groups, evaluation of existing and integration of new improved measurement techniques into operational measurement systems, and exchange of data for cross-validation of studies. A more specific list includes:

1. Benoit Donnet presented his latest work on topology discovery using mrinfo [26], which inspired CAIDA's use of mrinfo for our topology mapping project.
2. Benoit Donnet is visiting Matthew Luckie in NZ in May/June 2010 to collaborate on cross-validation of mrinfo data using CAIDA's Ark infrastructure.
3. Benoit's mrinfo dataset is used by several people (Ethan Katz-Bassett, Justine Sherry, Mehmet Tozal, ...) for cross-validation of alias resolution inferences.
4. Rob Beverly presented his work on directed probing, which motivated Young Hyun of CAIDA to refine our on-demand topology probing API to facilitate Rob's use of Ark to study and improve his directed probing methods. Rob noted that the Ark platform provided an invaluable means to prototype and launch his directed probing experiments.
5. Young provided CAIDA topology data analysis results to Justine Sherry and Ethan Katz-Bassett (U. Washington) for comparison to their own iPlane alias resolution work.
6. Matthew Luckie (U. Waikato) and Amogh Dhamdhere (CAIDA/UCSD) are now collaborating on a study to quantify the extent of false links in traceroute-based IP topologies.
7. Ethan Katz-Bassett (UW) is using CAIDA's Marinda and Ruby tools to work with scamper to implement reverse traceroute in production. Matthew also added improved prober code to scamper to support Ethan's experiment.
8. Ethan Katz-Bassett (UW) is using the ISI hitlist through PREDICT, to fill gaps in their measurement target lists to support reverse traceroute.

9. RIPE is now using CAIDA's Ark dns-names data to feed into their Internet Number Resource DataBase (INRDB) and allow DNS results to be returned via a query interface to INRDB (REX).
10. Following the presentation by Eotvos university at AIMS-2, DIMES IP geolocation evaluation study now includes the Spotter tool, and the two institutions jointly published a technical report of the study [29].
11. iPlane data is now mirrored by RIPE to enable better sharing of iPlane's measurements of topology and link properties.
12. Nick Feamster (GaTech) is discussing a collaboration with Ethan (UW) on route-poisoning experiments.
13. RIPE NCC is looking into using Marinda/tuple space for active measurement infrastructure.
14. Peter Haga and Peter Matray (ETOMIC and Eotvos Lorand University) used the ground truth geolocation data provided by CAIDA, to validate their Spotter service [28]. They also collaborated with Christos Papadopoulos (Colorado State University) to get access to lists of blacklisted IP addresses to geolocate.
15. Dave Choffnes (Northwestern U.), who presented part of his thesis work at the workshop, is now an NSF CI postdoctoral fellow at U. Washington working with Ethan Katz-Bassett, Tom Anderson, and Arvind Krishnamurthy.
16. Benoit Donnet (Universite Catholique de Louvain) and Noa Zilberman (Tel Aviv) are both using Northwestern EdgeScope data.

4. UPDATES TO THE AIMS-1 RECOMMENDATIONS

The 1st AIMS workshop [18] developed a set of recommendations intended to advance the field of Internet active measurements. While these recommendations represent a multi-year "wish list" of Internet researchers, we offer a review of progress made on them in the last year. Tables 1, 2, 3 below summarize our assessments.

Figure 1 shows the results of the mind-mapping exercise during the AIMS 2010 Workshop brainstorming session.

ACKNOWLEDGMENTS. The workshop was supported by the National Science Foundation via grants CNS-0958547, OCI-0963073, and CNS- and by the U.S. Department of Homeland Security (DHS) Science and Technology (S&T) Directorate via contract N66001-08-C-2029. We thank all participants for their insights.

The research community needs to introduce and agree upon standards and best practices to promote a diverse and heterogeneous field of Internet active measurements.

	<i>Recommendation</i>	<i>Current status</i>
1	<i>develop standard APIs for measurement systems, standardize tool output, enable tool sharing on different platforms</i>	Work in progress. While no single authority leads or coordinates this activity, informal collaborations (partial list in Section 3) are yielding some convergence regarding efficiency of systems and formats.
2	<i>publicize the best available data, document them as ground truth, provide comprehensive statistical characterization, make these data easily downloadable</i>	Limited progress. e.g., CAIDA has contributed to this community priority via data sets such as the Internet topology data kits (ITDK) [15], and created another ITDK shortly after the workshop. RIPE's NCC Data Repository is another project aimed at improving data access for researchers. Researchers typically would like to share their data, but documentation and improving accessibility would require dedicated funding or other incentive to provide, further complicated by the required secrecy of most data-sharing arrangements. In other words, blocked on economic and policy issues.
3	<i>design flexible, easily extensible measurement infrastructure platforms capable of running various tools and types of measurements at Internet scale</i>	Researchers maintain, improve, and rearchitect measurement infrastructures based on experience, but the scope, availability, and usability of such infrastructure is typically resource-constrained.
4	<i>provision for continuity of measurements, dissemination of data, with long-term archiving of data to study historical trends</i>	This task deserves more attention from government, funding, and policy-making agencies. Longitudinal data collection and analysis continue to elude the research community, again based on resource constraints and relative priorities.
5	<i>maintain no-probe lists based on requests</i>	Researchers typically maintain their own <i>ad hoc</i> lists, but there is no mechanism in place to share/combine/standardize such lists across research groups. Combining lists would require some annotation complexity since different research questions need different types of measurements across potentially different populations of hosts, i.e., a single one-size-fits-all list would unnecessarily constrain the pool of hosts available for a given project.

Table 1: Recommendations on standards.

The lack of consistent guidelines for Internet measurement limits the recognized legitimacy and sustainability of Internet measurement systems.

	<i>Recommendation</i>	<i>Current status</i>
1	<i>replace obsolete RFC1262 with a more current document</i>	Not done, blocked on working group volunteering.
2	<i>create a report on ethical guidelines for Internet active measurement research</i>	One relevant thread: DHS's PREDICT [6] project has a working group attempting to codify ethical guidelines for Internet research, motivated by the need to advance cybersecurity research while respecting evolving expectations of privacy [14].
3	<i>facilitate interaction between Internet researchers in Institutional Review Boards (IRB) that overview and regulate human research activities at individual institutions [16]</i>	Several research groups have pursued and received approval from their IRBs for active measurement. In some cases the application is dismissed with the explanation that active Internet measurements do not involve human subjects. The ethics report mentioned above will further inform this politically sensitive subject.
4	<i>identify important research questions/problems in the field of Internet research where macroscopic active measurement can have a positive impact</i>	Workshops like AIMS provide a helpful forum for researchers to discuss priorities, benefits, and expected outcomes of active measurement research activities. Figure 1 depicts the connectivity of important concepts that underlie the research conversation about Internet topology.

Table 2: Recommendations on guidelines.

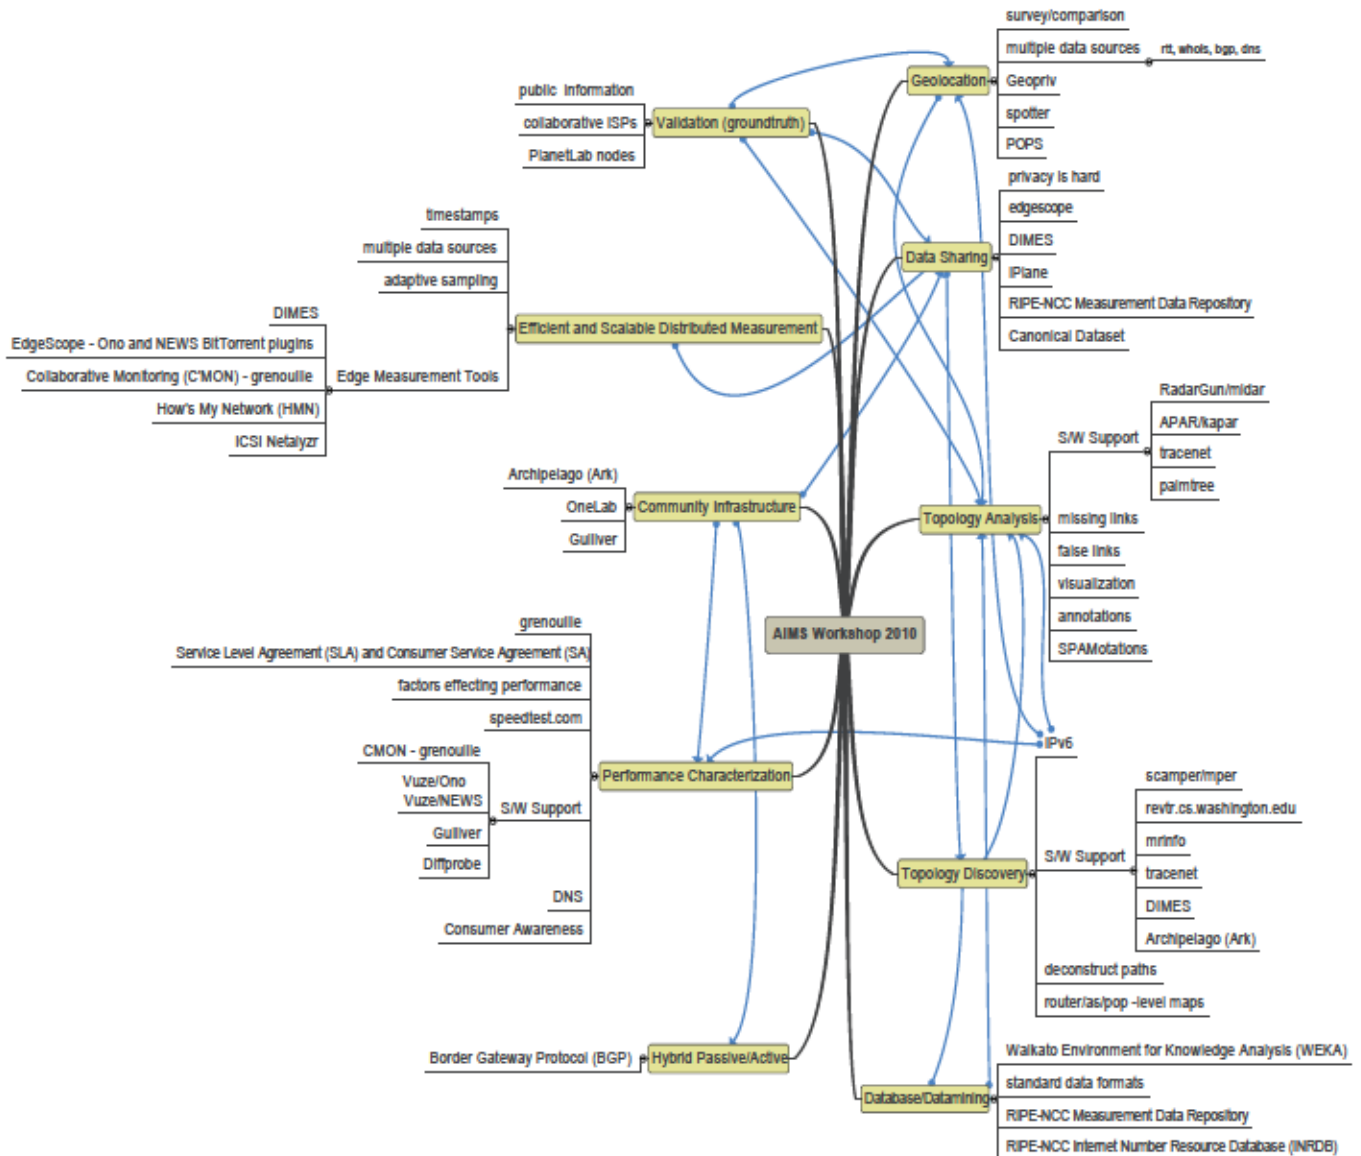


Figure 1: Topic map created during the AIMS Workshop brainstorming session.

The research community must increase transparency of Internet measurements and better communicate utility of results to broader communities affected by measurements (legal, political, operators, users).

	<i>Recommendation</i>	Current status
1	<i>create a central easily accessible database of planned or ongoing Internet experiments</i>	Not done. At present, this task can only be done as a volunteer effort, e.g., via a group-editable wiki page. Someone needs incentive (and/or funding) to pursue it.
2	<i>release source code for tools used for publications</i>	No available survey of how often this happens.
3	<i>consider other means of communication (i.e., blogs, mailing lists, automated announcements) to keep other communities informed of Internet measurement research experiments</i>	Some research groups do use blogs to announce experiments, as well as continued use of operational mailing lists like NANOG's and RIPE's.
4	<i>increase visibility and usability of data (including formatting standards [27]), relevance of data to users, and exposure of implications of studies based on data</i>	Some measurable progress, outlined in several of the workshop talks.
5	<i>inform debate on clean-slate Internet architecture</i>	Measuring IPv6 deployment and structure can inform debates on clean-slate Internet. Deeper discussion needs cross-fertilization among communities, which in turn requires prioritization by funding agencies.
6	<i>discuss with academics, operators, and funding agencies how many measurement infrastructures are needed, for what purposes, and if there are more effective ways of funding them</i>	Ongoing, worth noting considerable progress both at and subsequent to AIMS workshops.
7	<i>enable interaction and technology transfer between three main players in the field of Internet research: academic laboratories, commercial enterprises, and government institutions</i>	Still needs considerable help from policy side, with improved technology transfer methods including data-sharing models. [20]

Table 3: Recommendations on transparency.

5. REFERENCES

- [1] iPlane. <http://iplane.cs.washington.edu>.
- [2] AquaLab@Northwestern. EdgeScope – Sharing the view from a distributed Internet telescope. <http://www.aqualab.cs.northwestern.edu/projects/EdgeScope.html>.
- [3] Richard Barnes. Ietf working group on geographic location/privacy, 2009. <http://tools.ietf.org/wg/geopriv/>, rbarnes@bnn.com.
- [4] Robert Beverly, Arthur Berger, Young Hyun, and kc claffy. Understanding the Efficacy of Deployed Internet Source Address Validation Filtering. In *8th ACM SIGCOMM/USENIX Internet Measurement Conference (IMC)*, November 2009.
- [5] Bradley Huffaker and kc claffy. CAIDA’s Geolocation Tools Comparison, 2010. <http://www.caida.org/projects/cybersecurity/geolocation/>.
- [6] CAIDA. The DHS PREDICT project. <http://www.caida.org/projects/predict/>.
- [7] CAIDA. Netgeo ref. <http://www.caida.org/utilities/netgeo/>.
- [8] CAIDA. ISMA 2009 AIMS - Workshop on Active Internet Measurements, 2009. <http://www.caida.org/workshops/isma/0902/>.
- [9] CAIDA. ISMA 2010 AIMS-2 - Workshop on Active Internet Measurements, 2010. <http://www.caida.org/workshops/isma/1002/>.
- [10] Kai Chen, David Choffnes, Rahul Potharaju, Yan Chen, Fabian Bustamante, and Yao Zhao. Where the Sidewalk Ends: Extending the Internet AS Graph Using Traceroutes From P2P Users. Rome, Italy, December 2009.
- [11] David Choffnes and Fabian Bustamante. Pitfalls for testbed evaluations of Internet systems. *SIGCOMM Computing Communications Review*, 40(2), April 2010.
- [12] David R. Choffnes and Fabián E. Bustamante. Taming the torrent: A practical approach to reducing cross-ISP traffic in peer-to-peer systems. In *Proc. of ACM SIGCOMM*, Seattle, WA, Aug. 2008.
- [13] David R. Choffnes, Fabián E. Bustamante, and Zihui Ge. Crowdsourcing service-level network event detection. In *Proc. of ACM SIGCOMM*, New Delhi, India, August 2010.
- [14] Douglas Maughan Erin Kenneally, Michael Bailey. A Framework for Understanding and Applying Ethical Principles in Network and Security Research. In *Ethics in Computer Security Research (WECSR 2010)*, January 2010.
- [15] Cooperative Association for Internet Data Analysis. Internet Topology Data Kit, April 2010. <http://www.caida.org/data/active/internet-topology-data-kit/>.
- [16] Simson Garfinkel and Lorrie Faith Cranor. Institutional review boards and your research: A proposal for improving the review procedures for research projects that involve human subjects and their associated identifiable private information. 2010. <http://simson.net/clips/academic/2010.CACM.IRB.pdf>.
- [17] Bradley Huffaker, Amogh Dhamdhere, Marina Fomenkov, and kc claffy. Toward Topology Dualism: Improving the Accuracy of AS Annotations for Routers. In *PAM*, 2010.
- [18] kc claffy, M. Fomenkov, E. Katz-Bassett, R. Beverly, B.Cox, and M. Luckie. The Workshop on Active Internet Measurements (AIMS) Report. *Computer Communication Review*, 39(5), 2009.
- [19] kc claffy, Young Hyun, Ken Keys, Marina Fomenkov, and Dmitri Krioukov. Internet mapping: from art to science. In *CATCH*, 2009.
- [20] Erin Kenneally and kc claffy. Dialing privacy and utility: a proposed data-sharing framework to advance Internet research. *IEEE Security and Privacy*, 2010. http://www.caida.org/publications/papers/2010/dialing_privacy_utility/.
- [21] K. Keys. Ip alias resolution techniques. *Computer Communication Review*, January 2010.
- [22] D. Krioukov, F. Chung, kc claffy, M. Fomenkov, A. Vespignani, and W. Willinger. The Workshop on Internet Topology (WIT) Report. *Computer Communication Review*, 37(1), 2007.
- [23] S. Laki, P. Matray, P. Haga, I. Csabai, and G. Vattay. A Model Based Approach for Improving Router Geolocation. *Computer Networks*, 54:1490–1501, 2010.
- [24] Harsha V. Madhyastha, Tomas Isdal, Michael Piatek, Colin Dixon, Thomas Anderson, Arvind Krishnamurthy, and Arun Venkataramani. iPlane: An Information Plane for Distributed Services. In *OSDI*, 2006.
- [25] Harsha V. Madhyastha, Ethan Katz-Bassett, Thomas Anderson, Arvind Krishnamurthy, and Arun Venkataramani. iPlane Nano: Path Prediction for Peer-to-Peer Applications. In *NSDI*, 2009.
- [26] Pascal Mérindol, Virginie Van den Schrieck, Benoit Donnet, Olivier Bonaventure, and Jean-Jacques Pansiot. Quantifying ascs multiconnectivity using multicast information. In *IMC ’09: Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, 2009.
- [27] Saverio Niccolini, Sandra Tartarelli, Juergen Quittek, Thomas Dietz, and Martin Swany. Information model and XML data model for traceroute measurements. <http://www.ietf.org/rfc/rfc5388.txt>.
- [28] S. Laki and P. Matray and P. Haga and T. Sebok and I. Csabai and G. Vattay. Spotter: A Model Based Active Geolocation Service, 2010. in preparation.
- [29] Yuval Shavitt and Noa Zilberman. A study of geolocation databases. Technical report. <http://arxiv.org/abs/1005.5674>.
- [30] Fabien Viger, Brice Augustin, Xavier Cuvellier, Clémence Magnien, Matthieu Latapy, Timur Friedman, and Renata Teixeira. Detection, understanding, and prevention of traceroute measurement artifacts. *Computer Networks*, 52(5):998–1018, 2008.
- [31] Team Contact W3C Womer, Matt. Geolocation working group, 2009. <http://www.w3.org/2008/geolocation/>.