The 4th Workshop on Active Internet Measurements (AIMS-4) Report

kc claffy
CAIDA/UCSD
kc@caida.org

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 It has NOT been peer reviewed.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

On February 8-10, 2012, CAIDA hosted the fourth Workshop on Active Internet Measurements (AIMS-4) as part of our series of Internet Statistics and Metrics Analysis (ISMA) workshops. As with the previous three AIMS workshops, the goals 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 address future data needs of the network and security operations and research communities. This year we continued to focus on how measurement can illuminate two specific public policy concerns: IPv6 deployment and broadband performance. This report briefly describes topics discussed at this year's workshop. Slides and other materials related to the workshop are available at [1].

Categories and Subject Descriptors

C.2.3 [**Network operations**]: Network monitoring; C.2.5 [**Local and Wide-Area Networks**]: Internet; C.2.6 [**Internetworking**]: Standards; C.4.2 [**Performance of Systems**]: Measurement techniques—Active

Keywords

active measurement, Internet measurement techniques, validation

1. MOTIVATION

The AIMS workshop series was established to foster interdisciplinary conversation among researchers, operators, and government, focused on analysis of goals, means, and emerging issues in active Internet measurement projects. For four years, the workshop has helped stakeholders in Internet active measurement projects to communicate their interests and concerns, and explore cooperative approaches to maximizing the collective benefit of deployed infrastructure and gathered measurements. The final report from the first AIMS workshop [2] 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 and evaluate progress on this roadmap, build on previous achievements, refine our understanding of remaining problems and recognize new ones, modifying the course of progress as necessary.

The first workshop emphasized discussion of existing hardware and software platforms for macroscopic measurement and mapping of Internet properties, in particular those related to cybersecurity. The second workshop included more performance evaluation and data-sharing approaches. For the third workshop we expanded the agenda to include active measurement topics of more recent interest: broadband performance, gauging IPv6 deployment, and measurement activities in international research networks. This year we continued to focus on IPv6 deployment and broadband performance issues, and specifically how measurement in these areas can inform public policy debates. The workshop achieved its goals: sharing techniques and plans for IPv6 and broadband measurement; discussing issues inhibiting progress; and enabling collaboration, coordination, and data sharing among participants. This report briefly describes topics covered at this year's workshop. Slides and other materials related to the workshop are available at [1].

2. IPV6

Robert Kisteleki of RIPE NCC presented an update on their new active measurement infrastructure launched in November 2010 – RIPE Atlas. With 1200 probes (mostly in Europe, about half IPv6-capable) as of February 2012, RIPE NCC's goal is to scale up to several thousand, potentially tens of thousands, of vantage points and support simple built-in and user-specified measurements via an API (and eventually a web-based interface). The Atlas nodes are keychain-sized (8MB RAM, 16MB flash), running a reduced Linux distribution, Ethernet-connected to the Internet, USB-connected (and powered) and thus trivial to deploy. Atlas nodes perform a small set of measurements (ping and traceroute) to a provided list of targets, and limited DNS query measurements to root servers. As of February 2012, RIPE NCC was receiving a steady aggregated stream of about 10,000 Atlas-generated data points per minute, which they can plot in real-time to visualize metrics of Internet health. They are also developing a system to incent probe hosting by allowing hosting sites (including ISPs) to accumulate credits that they can use to request measurements from the system. Robert invited ideas for other measurements to integrate into Atlas.

George Michaelson gave an update on APNIC's measurements of IPv6 capability at the edge, which they undertake using two methods: Javascript-induced browser measurements triggered by a website script; and flash-induced browser measurements triggered by advertisement networks such as Google. The Javascript measurements require websites willing to run a script triggering a token IPv6 download, and the resulting data suggests some bias. They obtain about 50K-100K hits per day across 20-30 participant web sites running the Javascript, and 50K impressions (data points) per day as a return on a $20 daily investment into Google's advertising network. They stratify the results by originating AS to determine which networks are providing IPv6 support to browsing clients. In January 2012 across both sets of measurements they gathered data from over 20,000 unique IPv4 Autonomous Systems (or 50% of the 40,000 total IPv4 ASes), with about 1400 of them observably IPv6-capable. Of approximately 900,000 sample points per week, 0.3% of clients preferred IPv6, with minor variations and regional differences but no discernibly upward trend over the last year. They also observed substantial levels of "hop-over" behavior, where matching IPv6 and IPv4 prefixes originated from different ASes, presumably because the IPv4-transit AS does not yet support IPv6, a prominent problem with broad CPE deployments, e.g,. access providers. They noted that they have only scratched the surface in analyzing their data; they have not yet analyzed RTT, MTU/MSS, pMTU, and will continue to post results on their web site (http://labs.apnic.net/ipv6measurement).

Ann Cox presented preliminary work by Arthur Berger of Akamai (not present) comparing IPv4 and IPv6 path latencies between dual-stacked Akamai nameservers. He ran 9 months of continual ping measurements from three vantage points in San Jose CA, Dallas TX, and Reston, VA, to about 7,000 globally distributed dual-stack (non-Akamai) nameservers, recording approximately 44M measurements. Usually the IPv4 paths performed better, although roughly 25performed better. The presence of tunnels supporting IPv6 transport to a nameserver was correlated with higher latency and loss to that nameserver than native IPv6 transport. More surprisingly, IPv6 auto-tunnels (6to4, Teredo) were correlated with increased latency of native IPv4 transport to the same nameserver, largely due to differences in the continental distribution of 6to4 and Teredo, and legacy operating system issues at the endpoints.

Rob Beverly of the Naval Postgraduate School described a proposal and prototype instrumentation to use abusive IPv6 traffic to reveal properties of the IPv6 Internet. After demonstrating that approximately 1-2% of known abusive sites advertise AAAA records, he instrumented and operated an IPv6 spam honeypot and corresponding authoritative nameserver starting on World IPv6 Day on 8 June 2011. In one month, they received 14 spam email messages via IPv6, covering a variety of scams, languages, origin autonomous systems, and all apparently from server hosts rather than compromised user hosts (e.g. bots at the edge). Although the results were preliminary, they demonstrate that abusive IPv6 traffic does exist, and that, in the future, it may provide a valuable source of opportunistic IPv6 measurements.

Richard Barnes from BBN proposed some techniques for intelligent adaptive probing of the IPv6 address space. His proposed technique starts with advertised BGP prefixes, and then approximates a random scan across 16 subnets by adding 4 random bits to the selected prefix. He also exploits WHOIS information, since it is sometimes registered at a finer granularity than BGP. Finally, he used intermediate addresses gathered in previous scans to guess at a network's IPv6 addressing scheme. Each technique revealed some additional IPv6 topology.

We continued the next morning with IPv6 measurement talks, starting with Young Hyun of CAIDA describing and demonstrating CAIDA's new "topology-on-demand" real-time measurement infrastructure for supporting IPv4 and IPv6 topology measurements from multiple vantage points simultaneously. This service is currently implemented as a (Ruby-)scriptable interface for performing IPv4 and IPv6 traceroutes and pings from all 57 (as of February 2012, 28 with IPv6) Ark monitors. The on-demand measurement architecture uses the Marinda tuple space, which also supports the underlying Ark platform, for decentralized communication and coordination. One goal is to support varying levels of user sophistication and needs, from simple one-time probing to a set of pre-defined targets, to adaptive feedback-driven measurements in the language of the user's choice. CAIDA eventually plans to support a web-based interface, the ability to re-use previously requested measurements, as well as other services such BGP queries, mapping IPs to ASes, prefixes, routers, and geolocations.

Geoff Huston offered a comparison of the dynamics of IPv4 and IPv6 BGP routing planes. Publicly available IPv6 routing tables have grown from zero to 8,000 routes in the last six years. The number of IPv4 updates has remained at about 100K/day since 2008 (updating about 20K prefixes/day), while the number of IPv6 updates has risen from 1000 to 10,000 in the same period (and from 100 to 1,000 updated prefixes), or approximately the same growth rate as the IPv6 routing table. He also found that in IPv4, average convergence latency (measured in time or number of updates) and average AS path length have been mostly stable for years. It is harder to draw trends for the IPv6 network since there is wider variance among views, in terms of stability and convergence behavior. But about 10% of IPv6 prefixes observed in global routing tables exhibit instability properties, while since 2004 similarly unstable IPv4 prefixes have dropped from 10% to 5% of advertised prefixes. For IPv6 the average convergence time was similar to IPv4, around 100 seconds, until 2009 when IPv6 average convergence time and average number of updates required to converge increased over 5X and 30X, respectively, and has only recently started to drop. Geoff suspected that IPv6 tunnels (encapsulation and decapsulation code paths) may be part of the problem, as well as lack of necessary IPv6 peering expertise. Another methodological problem with comparing IPv4 and IPv6 trends of only the the last eight years is that IPv4 routes have slowed to linear growth for the last decade (at least), but IPv6 is still in an exponential growth phase; a more meaningful analysis would require comparing current IPv6 update dynamics with those of IPv4 back when its routing table was growing exponentially, or alternatively comparing prefix updates on a per AS-level rather than in aggregate.

Dan Massey of Colorado State University presented another view of IPv6 BGP dynamics, based on data from his 6watch project (6watch.net). He confirmed Geoff's view that different vantage points see substantially different views of IPv6, some of which is due to aggregation, testing, and simply different prefixes being observed at different locations. Many missing prefixes are covered by an aggregate, presumably mitigating the risk of packet loss. Half of the unreachable prefixes are transient for a day or two and half are persistently unreachable. Most (87%) of the 4,000 IPv6 prefixes in the Route Views table at the beginning of 2011 persisted for the whole year, but since there was considerable growth in 2011 and 2012, there are still several thousand noisy IPv6 prefixes. Geoff suggested that maybe some providers are limiting the length of IPv6 prefix they carry, as Sean Doran famously did decades ago with IPv4 /20 prefixes. He also observed that the lack of officially private space in IPv6 leaves people sometimes using public IPv6 addresses as private, which then leak, exacerbated by (again) lack of IPv6 clue and support in network management devices. Dan did not however conclude from the data that people are trying IPv6 and turning it off, and posed the question: what do we want to measure about a new prefix after it is first seen? He hoped for a way to measure the actual impact of churn on traffic, by combining data plane and control plane measurements.

Emile Aben, beaming in from Amsterdam, presented recent work on partial reachability of IPv4 vs. IPv6 from the data plane perspective. He performed six iterations of full mesh ping probes from 23 Ark nodes with usable IPv6 connectivity to about 1273 dual-stacked destinations from the Alexa 1M, spread across 775 IPv4 ASes and 716 IPv6 ASes. He found 1.2% of the IPv4 ASes experienced partial reachability, i.e., some Ark node was unable to reach a destination in the AS. The corresponding number for IPv6 ASes was 9%, or 64 of the 716 IPv6 ASes. He then analyzed measurements from Atlas nodes at the edge to four fixed destinations: RIPE, the d-root nameserver, and two content distribution web sites. He managed to get an IPv6 connectivity problem to d-root (inadvertent filtering of a /48 on ARIN's critical infrastructure list) fixed by showing them the measurements. Using 15 days of Atlas measurements, he observed that 6to4 connectivity tended to have 5X (10X) as many (temporary) partial reachability issues as native IPv6.

Matthew Luckie of CAIDA closed the IPv6 half of the workshop with some trends and tidbits of IPv6 deployment based on 4,800 dual-stacked ASes from seven ASes providing both IPv4 and IPv6 BGP feeds to Route Views. He explored several hypotheses related to the maturity of the IPv6 Internet. The first hypothesis was that IPv4 and IPv6 level paths should grow to be more congruent as the IPv6 network matures, i.e., the edit distance between the IPv4 and IPv6 AS paths to the same destination should be moving toward zero for a larger fraction of dual-stacked paths. BGP data for the seven analyzed peers in Route Views did show a weak trend in this direction, although the data is biased by Hurricane Electric's provision of IPv6 transit in many paths where it is not in the IPv4 path. Removing Hurricane Electric as an outlier would yield an interesting "what-if" but there is no sound methodology for doing so using available BGP data. Matthew next used a snapshot of data (30 January 2012) from PeeringDB (www.peeringdb.net) to explore whether IPv6 capability correlated with network business type, reported bandwidth capability, geographic region, or RIR exhaustion in the network's region. Of the 2,622 ASes represented in PeeringDB for this snapshot, 60% advertised themselves as IPv6 capable. Unsurprisingly, and consistent with previous reports, more transit providers ("NSPs" in PeeringDB) than access, enterprise, or content providers reported IPv6 capability. The higher the reported traffic volume of the network, the more likely it reported IPv6 capability. Matthew also used Ark measurements to study performance differences between IPv4 and IPv6 path RTTs. Of the 20 Ark machines able to provide statistically significant samples of dual-stacked ASes, 16 (80%) of them observed larger IPv6 RTTs to dual-stacked ASes. These observations are mostly consistent with Akamai's 2010 data presented earlier, which found 75% of observed nameservers had larger IPv6 than IPv4 response times. (Emile's analysis of paths during 2011 World IPv6 Day showed only 60% of path RTTs slower for IPv6 [3].)

While there was significant similarity across different types of IPv6 measurement presented (data plane performance was a theme), the results were not always consistent, nor did they tell the complete story. Efficient and scalable measurement of the characteristics of the consumer edge is particularly relevant for gauging IPv6 deployment, but several attendees expressed interest in richer performance evaluation of IPv6 relative to IPv4, with a particular aim to improve IPv6 performance issues.

We did a roundtable survey at the end of the first day, and revisited it the end of this last IPv6 talk, asking participants to predict which of the following statements would best characterize the nature of IP transit in ten years (number of votes received in parentheses): (a) IPv6 is mostly deployed (20); (b) We're still struggling along trying to get IPv6 deployed much like we are today (4); (c) IPv6 is mostly dead and CGNs are everywhere (11); and (d) Something beyond IP has taken over (1 wishful thinker). The survey inspired a lively discussion of forcing functions for IPv6 deployment, in particular content versus carriage and which has more incentive and capital to push IPv6 deployment. Geoff predicted that content was taking on such an increasing proportion of the communication industry's capital that the content industry would force IPv6 on carriers one way or another. He challenged us to estimate how much we each personally spend on bit transport (carriage) per month (cable+phones+data plans) versus content, acknowledging the caveat that basic cable television was in both categories so needed special treatment. The results of this content-vs-carriage survey were eye-opening, revealing that most people were spending about 2:1 on carriage:content, but the trend was rapidly inverting this ratio, i.e., moving toward spending more on content than carriage. Geoff and a few other participants were ahead of this curve, spending at least five times (5X) as much on content as on carriage, and believed this trend would inevitably continue as content options proliferate. The content industry certainly has a strong incentive to not live behind CGN'ed walled gardens, and may well accumulate sufficient capital leverage to compete, lobby, or otherwise invest their way out of it.

3. BROADBAND PERFORMANCE

Beginning after lunch on Day 2, we began the second theme of the workshop: broadband performance measurements, a topic of expanding interest to consumers, government, and industry. This topic was more narrowly focused, in and some ways more mature than the IPv6 topic since IPv6 technology itself is still maturing, as is IPv6 measurement technology, and there are many dimensions beyond performance that merit measurement.

Tiziana Refice of Google kicked off the session with a status update on M-Lab (Measurement Lab), the collaborative server infrastructure founded by Google to support performance measurements. Each M-Lab node (76 servers in 22 locations as of February 2012[1]) is a PlanetLab-like platform with Web100 instrumentation to support performance diagnostic output. M-Lab hosts several open-source software tools, including NDT from Internet2 (by far the most heavily used on M-Lab), and Pathload2 and ShaperProbe from Georgia Tech. M-Lab also supports measurements from two Router-based tools, SamKnows and BISmark, both of which were presented later in the workshop. M-Lab has accumulated 460 Terabytes of data so far, running about 150,000 tests per day. Regulators in Greece and the U.S. have used M-Lab data in constructing national broadband maps. Google offers research grants to port new tools to M-Lab.

Sam Crawford of SamKnows gave some background on their measurements of actual vs. advertised end-to-end consumer broadband performance, which they have now conducted for four regulators in Britain, the U.S., the E.U., and Singapore. They use hardware probes with a customized Linux stack to measure end-to-end latency, DNS failure rates, and other metrics. For the first phase of the FCC-funded study, they used 7000 hardware probes to study the actual performance of the top 16 fixed-line ISPs in the U.S., representing 85% of consumers. They used the nearest M-Lab server to the consumer probe for off-net measurements, and servers inside the monitored ISP's network for on-net measurements. They found a surprisingly small difference between the off-net (M-lab) and on-net (ISP) measurements, about 0.4% for the first week of February 2012. They ran the U.S. experiment for several months before involving the ISPs to avoid active interference to influence results, and only ran measurements when the home network was idle. Tests ran between once and 600 times per hour depending on the variance in the results. They found mostly good news, that no provider was delivering substantially less than they advertised, and that many cable service tiers exceeded 100% of their advertised upstream rate (Partha later questioned whether this result was due to the 30s TCP stream used to measure the sustained upstream rate, which will yield misleading results if traffic shaping kicks in after 25s, which it does for several cable providers including Comcast.) For the second (future) phase of the FCC experiments SamKnows will be using OpenWRT-based hardware probes with IPv6 support (although M-Lab servers do not support IPv6 yet so the future of IPv6 support is unclear), and supporting additional tests including loss-under-load. SamKnows also hopes to expand into mobile broadband measurement.

Ahmed Elmokashfi of Simula Research Laboratory (Norway) gave a short summary of a project measuring mobile broadband performance of multiple operators in Norway, motivated by the government's interest in ensuring network availability during elections (to get voter registration verification data reliably from the voting centers to the central location, not to vote electronically). They ran ping measurements every second over both fixed and mobile lines. They labeled a gap of > 15 sec as downtime, and found high variance in downtimes across operators, up to as high as 10 minutes per day (Telenor). Most downtime episodes lasted less than 3 minutes, and failures tended to cluster in time and space: 70% of the failures were within 100 sec of each other, and geographically closer monitors were more likely to fail simultaneously. The three mobile providers exhibited availability of 99.51%, 99.83%, and 99.73% during the observed interval but combining the three operators increased the overall availability to 99.98%. Although this measurement project was terminated after the election, Simula is planning a new infrastructure based on a micronode model, e.g., BeagleBone single-board computers.

Christian Kreibich from ICSI introduced a new project from his group – Fathom, a (Firefox) browser-based network measurement platform to support user measurements. After motivating and discussing a number of potential implementation paths for building a browser-based measurement platform, he reviewed the architecture and security model of Fathom, followed by a demonstration. Although only in prototype stages, they are eager to build a community of users and open source contributors (http://fathom.icsi.berkeley.edu). There was follow-up discussion of why they chose Firefox as the base platform (It supported socket access via Javascript.) and how hard it would be to port

[1] http://measurementlab.net/mlab_sites.

Fathom to other browser clients. (It varies: Safari would require a plug-in, Internet Explorer will soon not allow extensions at all, Chrome is doable in native code rather than a plug-in.)

Nick Weaver gave a status update on the next version of Netalyzer (to be released soon at http://netalyzer.icsi.berkeley.edu). In addition to increasing the speed of measurement execution, they have added enhancements to improve the GUI and facilitate embedding the tool in other code. They have added JSON support, which among other features enables one to generate individual result URLs and share them. The data-sharing model is still challenging; for specific questions they can perform custom analyses in their database, or extract specific fields for a class of sessions. But they cannot share the whole dataset since some tests reveal outbound security postures and other information (e.g., potentially vulnerable caches or proxies, software version information, particulars of user NATs that may reveal vulnerabilities).

For the final talk of Day 2 we had a taste of something different. Casey Deccio from Sandia Labs led a demo and discussion of DNSViz, an active DNS measurement and visualization tool he created to provide insight and guidance to improving the quality of DNSSEC deployment. DNSSEC's potential to dramatically affect residential broadband performance was loudly demonstrated in January 2012 when NASA.gov had a DNSSEC signing error[2] that prevented Comcast's DNSSEC-capable residential customers (i.e., all of them) from being able to reach NASA.gov. Casey showed how DNSViz could actively illuminate the problematic NASA domain from a given vantage point, and make results available for visual analysis on http://dnsviz.net/ (which he did). He has also used DNSViz to perform a DNSSEC deployment survey, polling 2700 production signed zones from May 2010 - July 2011, identifying and classifying over a thousand misconfigurations (almost half of the signed zones!), including IPv6 inconsistencies. His disconcerting conclusions include that administrators are not detecting and correcting their own DNSSEC problems in timely fashion, nor are they tending to roll their keys, and resolver operators are learning about third-party misconfigurations from their own disgruntled customers. Casey's future plans for DNSViz include augmenting the analysis with passive measurements, smarter probing, alerts when misconfigurations are detected, and a RESTful API to support third-party monitoring tool development.

Day 3 began with our second remote talk, by Jim Gettys from Lucent's Bell Labs, on the topic of bufferbloat (bufferbloat.net), and what we can do about it. He began with a demo of web browsing latency with and without a large file copy operating in parallel. The large file transfer increased the interactive latency by a factor of 15, and reducing the size of the buffer in the home router increased performance ten-fold. The problem is simple: the buffers in the home router fill up with the file transfer payload, and are unavailable for the interactive traffic. But the result is a broken network edge for the most common (pervasive) home broadband scenarios, i.e., those mixing bulk transfer and interactive traffic without prioritization. Gettys emphasized that what is really needed is active queue management (AQM) such as RED (Random Early Detection/Drop). But appropriately tuning RED buffers requires expertise that cannot be expected from typical users, so marketing forces tend to lead to architectures with bigger buffers to minimize packet drops, even though packet drops are essential to TCPs congestion control mechanisms. Geoff pointed out that we already know that TCP's feedback algorithm is ill-suited to fair sharing across a wide range of flow demands, which is why researchers developed algorithms such as FAST TCP to use RTT or queuing delay rather than packet loss as an indicator of full buffers (i.e., congestion signal). (FreeBSD9 has an RTT-sensitive TCP so this hypothesis is somewhat testable.)

Renata Teixeira from CNRS and UPMC Sorbonne University updated us on her group's work using UPnP for the challenging problem of identifying home network problems from measurements. They explored the utility of the *universal plug and play* (UPnP) feature which, among other features, allows end hosts to query certain home gateways for traffic counters (e.g., bytes and packets in both directions) and other meta-data about the router's capabilities. They found that UPnP was supported in about 35% of the 100K home gateways they probed. Although there were problems with the accuracy of query results, when answers were correct they useful information, including synchronization rates vs. measured bandwidth, cross-traffic, packet loss and traffic counters within the home versus to/from the Internet. Sam pointed out that in UK, the carriers are removing UPnP support for security reasons, but Renata had not observed a strong trend either way, though she acknowledged security and privacy issues. Sam also observed that SNMP counters would work too, if home devices supported SNMP.

Aaron Schulman of the University of Maryland reviewed the research he presented at IMC 2011, assessing the quality of residential broadband performance during weather events, including lightning, rain, wind, hurricanes, and tornados. He extracted 100M residential IP addresses to ping using reverse DNS mapping, and geolocated the IP address so he could correlate these locations with severe weather reports from the National Weather Service. He executed ping measurements from 10 locations in the U.S. to 100 homes in a weather alert area, every 11 minutes to avoid abuse reports from ISPs. He found strong correlations: double the number of failures during rain and quadruple in thunderstorms relative to clear conditions, perhaps due to power failures. Although it only works based on U.S. weather information now, Aaron wants to extend it outside the U.S. and incorporate traceroute data in addition to pings. The current bottleneck in his system is getting the weather information; Daniel offered to help get more efficient access to (pilot-accessible) real-time weather information for Europe.

Srikanth Sundesaran and Steve Woodrow from Georgia Tech brought us back to discussions of measuring performance bottlenecks on the web. Srikanth first discussed some analysis of FCC Samknows data from 5500 homes to 11 ISPs, in which he found that browser load time flatlined for bandwidth throughput over 8Mb/sec, suggesting that pages have many small objects and TCP is stuck in slow-start most of the time trying to download them. They also used their BISmark testbed (27 homes) to investigate the potential benefit of performance acceleration techniques within home networks, e.g., DNS caching, TCP connection caching, content caching. They found that all three types of caching help, especially content caching, and the level of performance improvement varies with site and usage characteristics. Geoff pointed out that Firefox is already offering many tuning knobs and optimizations, including DNS caching.

Steve Woodrow provided an update on the BISmark active broadband measurement infrastructure, which had 55 routers as of February 2012 deployed across North America, Africa, Asia, and Europe. They hope to ship up to 200 routers, mostly outside North America. By April 2011 they will have migrated their performance measurements to use M-Lab servers to get more accurate throughput. At that point the measurement data will be made public within a year or after the first publication about it, but they are also interested in collaborating with others who might be interested in accessing the data already collected, or running their own measurement experiments on BISmark.

Most of these measurements are geared toward rough assessments of performance rather than rigorously quantifying the service provided. One concern expressed by participants was that most active probing is too infrequent (whether per IP, AS, or subnet) to detect many failures; the interval would need to be on the order of minutes not hours.

Partha Kanuparthy from Georgia Tech gave two updates on performance measurement tools he co-developed: WLAN-probe, for diagnosing home wifi network performance problems; and Pythia, for high-bandwidth network measurements such as on Internet2 and DOE's backbones. WLAN-probe (currently a set of proof-of-concept scripts rather than a downloadable tool) distinguishes among three wireless performance pathologies common in home networks: congestion, low signal strength, and hidden terminals. He also described the state of Pythia, a tool developed for DOE with three objectives: detection, localization, and diagnosis of wide-area performance problems. For detection of problems,

[2] http://www.dnssec.comcast.net/DNSSEC_Validation_Failure_NASAGOV_20120118_FINAL.pdf

they flag deviations from baselines of delay, loss, and reordering. The next step is to diagnose the root cause of an outage, e.g., nature of congestion, routing, type of reordering, end-host effects. Each category could have a variety of root causes. For instance, congestion could be due to a persistent queue build-up, bursty traffic causing delay jitter, or too small (or to large) buffers. For next steps, they hope to add support for detecting additional performance problems, complete an open-source implementation and front-end for operators, and deploy and validate its accuracy on operational backbones.

Steve Bauer of MIT gave a talk on the State of TCP, based on data from M-Lab server side Web100 TCP state variables measured from over half a billion connections to M-Lab (175M unique client IPs) during NDT, Glasnost, SamKnows, and BISmark tests. He found that NDT tests tended to measure slower speeds (even from the same IP address), in part because 34% of the tests do not fill the pipe enough to generate any congestion signal, i.e., are performance-limited by the TCP receive window setting. This receive-window limitation setting is not just a legacy system problem, since many low-memory systems will default to a low receive window, impairing performance. One conclusion, echoing Jim's earlier talk, is that TCP tuning is important not just for optimizing performance, but is even essential to capturing valid performance measurements. Proliferation of mobile devices will exacerbate this sensitivity.

We added to the final session lightning talks and demos by anyone interested in presenting. John Otto proposed *namehelp*, a tool for improving the performance of content delivery networks (CDNs) when users are relying on remote DNS servers, which ''trick" the CDN into sending the user to a server nearby the remote DNS server rather than nearby the user. Namehelp "untricks" the CDN by (re)co-locating the client and resolver, acting as a transparent DNS proxy. Finally, Geoff and George of APNIC showed several cool videos of botnet scanning effects on heatmaps.

4. RESULTING COLLABORATIONS

The format, schedule and cross-section of participants promoted collaborations following the workshop, including:

1. Amund Kvalbein, Tiziana Refice, and George Michaelson subsequently presented their work at RIPE 64
2. BBN and CAIDA have had some follow-up discussion related to smart scanning for IPv6
3. Georgia Tech's BISmark and ICSI's Fathom project are discussing client measurement APIs that could allow code reuse from Fathom to BISmark and vice versa. ICSI has now received a BISmark router, is porting Bro to it, and considering putting a web server on it to support Fathom-driven experiments.
4. BISmark developers were also planning to work with Steve Bauer (MIT) to add web100/web10g instrumentation to the BISmark router, to support capture of detailed TCP statistics from both sides of a BISmark active measurement.
5. RIPE and CAIDA are co-authoring a paper on the state of IPv6 evolution for IMC this year.
6. Casey Deccio and CAIDA will try to use CAIDA's DNS traffic data to gain insight into DNSSEC deployment.
7. Srikanth (GaTeach) provided Aaron's project with BISmark observations of residential IP changes during failures. With RIPE Atlas data and Robert's help, Aaron Schulman examined weather-related failures observed from inside homes (allowing him to identify power vs. link failures) and outside the U.S.
8. Simula hopes to collaborate with Steve Bauer on metrics and tools to assess mobile broadband robustness.

Another collaboration that developed from the two previous AIMS workshops has resulted in a SIGCOMM 2012 paper called "LIFEGUARD: Practical Repair of Persistent Route Failures", by Ethan Katz-Bassett, Colin Scott, Italo Cunha, David Choffnes, Vytautas Valancius, Nick Feamster, Harsha V. Madhyastha, Thomas Anderson, and Arvind Krishnamurthy, most of whom have attended previous AIMS workshops.

5. WORKSHOP FEEDBACK

We did a final roundtable to solicit feedback on the workshop and recommendations for how to improve future workshops, including what topics participants wanted to see covered next year. (Although we did not revisit the recommendations developed at AIMS-1 and tracked for the last 2 years, we plan to revisit and potentially update them next year.) Participants appreciated the opportunity for dedicated focus on the two topics – IPv6 and broadband measurement – given their increasing prominence in public policy debates and daunting dearth of empirical data to inform these debates. Our knowledge gap is exacerbated by the lack of funding in the public policy sector to support scientific research on issues under their purview. With respect to IPv6, we should know more after this year's "World IPv6 Launch" (6 June 2012) which has more ambitious goals (i.e., leaving IPv6 enabled after the day is over) than last year's event.

With respect to broadband, some participants emphasized the value of dedicated infrastructure to support broadband access active measurement, to save people from building redundant tools and platforms. Others reminded us that in the U.S. the regulatory definitions of broadband are overdue for revision, and we are at early stages of determining what metrics should be weighted more in "scoring" consumer broadband connections to capture the end-user quality of experience (not just network performance).

During solicitation of topics of interest for next year's workshop, many participants wanted to focus on new topics, especially the the robustness and quality of mobile network performance measurement and evaluation, and measurements that can support engineering and management of emerging content-centric networks. Other topics of interest included active tools and methods to identify and even potentially detect large-scale outages either due to geopolitical or geophysical events, or to support more precise size and growth estimates of Internet infrastructure, adoption, and usage.

Our friendly funding agency representative reminded us of the need for better ways to transfer measurement technology from research to operations, particularly the private and government sectors. Several participants had suggestions for items to require of speakers next year: a slide on what the speaker wants to get out of the workshop; a slide on how data is or will be shared; data visualization. Another repeated request was for solicitation of work analyzing and correlating multiple types of data, e.g., active and passive.

ACKNOWLEDGMENTS. The workshop was supported by the National Science Foundation via grants CNS-0958547 and OCI-0963073, and by the U.S. Department of Homeland Security (DHS) Science and Technology (S&T) Directorate via contracts N66001-08-C-2029 and NBCHC070133. We thank all participants for their insights, and many thanks to Evi Nemeth for taking excellent notes at the meeting.

6. REFERENCES

[1] CAIDA. ISMA 2012 AIMS-4 - Workshop on Active Internet Measurements, 2012. `http://www.caida.org/workshops/isma/1202/`.

[2] 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.

[3] Emile Aben. Measuring World IPv6 Day - Comparing IPv4 and IPv6 Performance, August 2012. `https://labs.ripe.net/Members/emileaben/measuring-world-ipv6-day-comparing-ipv4-and-ipv6-performance`.