

The 5th Workshop on Active Internet Measurements (AIMS-5) 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 It has NOT been peer reviewed.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 6-8, 2013, CAIDA hosted the fifth Workshop on Active Internet Measurements (AIMS-5) as part of our series of Internet Statistics and Metrics Analysis (ISMA) workshops. As with previous 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. The workshop focus this year was on creating, managing, and analyzing annotations of large longitudinal active Internet measurement data sets. Due to popular demand, we also dedicated half a day to large-scale active measurement (performance/topology) from mobile/cellular devices. This report describes topics discussed at this year's workshop. Materials related to the workshop are available at <http://www.caida.org/workshops/>.

Categories and Subject Descriptors

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

Keywords

active measurement, Internet measurement, validation

1. MOTIVATION

For five years, the AIMS workshops have 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 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 and fourth workshops we expanded the scope to include broadband performance and IPv6 deployment measurements, including how such measurements can inform policy debates. This year the theme was creating, managing, and analyzing annotations of large longitudinal active Internet measurement data sets (terabytes over several years). Much of our content and format of the workshop

this year derived directly from feedback at last year's AIMS workshop, in particular the focus on active mobile network measurement, and outage detection and analysis. Feedback from last year also inspired us to request that each talk have slides on data sharing and visualization methods. This report summarizes topics discussed at this year's workshop. Slides and abstracts are available at <http://www.caida.org/workshops>.

2. MEASUREMENT SYSTEMS

Ann Cox (Dept. of Homeland Security) introduced the Internet Measurement and Attack Modeling effort funded by DHS's cybersecurity program, which emphasizes technology transfer of applied research and development to the private as well as government sectors. Tech transfer is a challenge with Internet measurement research, which requires substantial and sustained investment in operational measurement infrastructure. DHS encourages open source software development as an important vehicle for technology transfer. Related to the meta-data theme of the workshop, DHS also supports PREDICT, an Internet security-related metadata repository, where the data stays with its owners, and legal disclosure control documents frame appropriate sharing policies.

James Grace (FIU) spoke as a network operator for AmLight, an NSF-funded International Research Network Connection project. AmLight connects the U.S. to South America via two 10-GB circuits. They are also funded by NSF to work with CAIDA on traffic measurement, and they are adopting CAIDA's CoralReef package to support their traffic reporting needs, specifically flow reporting at VLAN as well as AS granularities. James expressed interest in sharing data, or meta-data, with researchers.

Bradley Huffaker (CAIDA/UCSD) introduced a new lightweight DatCat (www.datcat.org) for storing meta-data of Internet data used for research. The original DatCat architecture was highly customizable, which made it too complex for most people to use, so CAIDA is launching a new simplified design and user interface this year. Robert Kisteleki (RIPE NCC) suggested an API to facilitate pushing measurement metadata to DatCat, or alternatively, DatCat code to automatically pull measurement metadata from other databases.

Robert Kisteleki (RIPE NCC) described several monitoring systems operated by RIPE that result in large data volumes, the most popular of which is the Atlas project. Atlas consists of over 2500 active small cigarette-pack-sized USB-attached probes distributed to users around the world, supplying 60M data points per day, mostly pings to root name servers. He noted that about 17% of the probes they hand out are never turned on, and another 17% are down at any given time. The infrastructure supports some user-defined measurements. The biggest technical challenge is not in performing the measurements, but in collecting the data and organizing it in a way that makes it easy to use later. Their current bot-

tleneck is getting people to request measurements from the system. The probes are running at approximately 1% of their capacity, and capable of doing ping, traceroute, DNS queries, and SSL. About 5000 measurement requests had happened by February 2013, from approximately 1000 unique users around the world. They recently introduced an API to support some querying, and will soon support user-configured periodic measurements. They are still working on a data sharing policy that clarifies what part of the data set is considered public.

Srikanth Sundaresan (Georgia Tech) gave an update on the BISmark project (projectbismark.net), which is trying to measure the Internet from home networks. They have sent out to end users over 300 Netgear routers (about 130 were deployed as of February 2013) using OpenWRT images, and including several open source modules that measure performance parameters including upload and download bandwidth, packet loss rate, and jitter. BISmark nodes use Google's M-Lab servers as destinations for active measurements. A control server at Georgia Tech supports device configuration and remote access, and coordinates measurements. Some of the active measurement data is made available to researchers via M-Lab. (They are doing passive measurements too, but have not pursued IRB approval to share that data.) They want to support research in the community, including customizing measurement tools and/or co-hosting regional deployments for specific experiments. BISmark data supports several research projects, including Srikanth's own research on home wireless performance issues and Sarthak Grover's (Georgia Tech) study of end-to-end routing behavior and path dynamics. They are actively seeking collaborators, as well as home users to deploy BISmark nodes in under-represented areas.

Nick Weaver (ICSI) gave an update on the Netalyzer project (netalyzer.icsi.berkeley.edu), including details of the growing data set, IPv6 nuggets learned from the data, and tips on using Netalyzer in research projects. As of February 2013, they had accumulated 790K sessions from 530K unique IPv4 addresses, resulting in 180GB of raw data. They have begun limited data releases to researchers with specific questions that can be answered by extracting non-sensitive portions of the database. They will add new tests on request if feasible. Of the Netalyzer sessions they have analyzed, 23K of the IPv4 addresses (5% of the IP addresses in their study) can fetch data using IPv6. The tests reveal more broken fragmentation behavior in IPv6 than is observed in IPv4: many IP addresses either cannot send or cannot receive (or both) fragmented IPv6 traffic. Netalyzer now supports a json API, and functionality that allows people to see and share the results of their measurements via a common URL. They are developing a test suite to study DNSSEC deployment, including client transport and client-side validation.

Mehmet Gunes (U. Nevada) presented an update on the Cheleby Internet topology mapping system. They used Cheleby to collect traceroutes from PlanetLab hosts to many destinations, and try to infer subnet structure from the traces. One goal is to generate a synthetic topology and subnet structure that matches a realistic (i.e., observed) topology, which requires determining parameters of the graph that induce the underlying topology. They share data at <http://cheleby.cse.unr.edu>.

Dan Massey (Colorado State) talked about his BGPmon project (<http://bgpmon.netsec.colostate.edu>) and in particular how to support the holy grail of Internet topology measurement: effectively combining active traceroute and passive BGP data, at scale. His group recently re-architected the BGPmon platform, (which relies on RouteViews to get data from ISP's) to support additional peers, and solicited input for where future BGP data collectors should be located. Among other enhancements, BGPmon now sup-

ports XML and ASCII output stream formats. Matthew Luckie (CAIDA/UCSD) requested support for sending raw MRT data over a socket, and offered to send a patch.

Alistair King (CAIDA/UCSD) closed this session with an update on CAIDA's efforts to visualize and annotate huge volumes (4.5TB/month, compressed, over 150TB total, not all on spinning disk) of the unsolicited background traffic observed by the UCSD telescope. These data volumes prohibit real-time insight without massive data reduction techniques. CAIDA wants to use the data to detect outages or other macroscopic performance changes in different regions, to support an interactive query interface data, and to scale the system to millions of metrics and years of data. CAIDA's previous approaches are not capable of this scale, so Alistair led the design and implementation of a new platform, Corsaro, available on CAIDA's web site. He uses MaxMind for IP geolocation, Whisper for data aggregation, ZFS for file system support, and graphite for visualization. Robert recommended tsdb as another open-source database alternative for distributed, scalable time-series processing.

3. TOPOLOGY ANNOTATIONS

Sandor Laki (Eotvos Lorand U.) presented early results from analyzing the spatial structure of the Internet topology. They captured 400K traceroutes from 300 PlanetLab sites. They used the Spotter tool to geolocate the resulting 13K IP addresses (44K links). They characterized the links in the graph to try to determine which links were most frequently used in paths, and which cities were most diversely interconnected. Their graph did not take into account the underlying cable structure, only the inferred location of routers. For visualization, they used *igraph* for R, Google Maps, and Quantum GIS. They share their data via SQL queries to their network measurement virtual observatory.

Riad Mazloum (UPMC Sorbonne) talked about problems his team has encountered with understanding AS relationship annotations in inferred data, that seem inconsistent with rational economic behavior. Of the ASes using multiple-exit routing, a significant fraction (30customer, peering, and provider links at the same time, based on traceroute and BGP data. He solicited suggestions for explanations of the mismatch between the (CAIDA-)inferred relationships and observed paths. Phillipa Gill asked if Riad had correlated the BGP data with the source of the traceroute (which he had not), and suggested that hot potato routing behavior might explain his observations. Learning economic motivations would require communicating with the operators.

Matthew Luckie (CAIDA/UCSD) gave an update on CAIDA's AS Rank project (<http://as-rank.caida.org>), which he has led since January 2013. He developed a new algorithm to infer AS relationships using BGP paths, which unlike previous approaches, does not seek to maximize the number of valley-free paths, instead relying on three assumptions about the Internet's interdomain structure: (1) an AS enters into a provider relationship to become globally reachable; and (2) there exists a peering *clique* of ASes at the top of the hierarchy, and (3) there is no cycle of p2c links. He assembled the largest source of validation data for AS-relationship inferences to date, validating 34.7% of the 125K c2p and p2p inferences to be 99.6% and 98.4% accurate, respectively. His three sources of validation data included assertions from operators directly to CAIDA, policies encoded in RIPE's WHOIS database, and community strings that imply certain routing policies. Using these inferred relationships, he evaluated three algorithms for inferring each AS's *customer cone*, defined as the set of ASes an AS can reach using customer links. He presented graphs of customer cone size over time that depict the rise and fall of large transit providers over the last fifteen years, including recent claims about the flatten-

ing of the AS-level topology and the decreasing influence of tier-1 ASes on the global Internet.

Ang Chen (U. Penn.) presented techniques for efficiently analyzing massive traceroute data to support topology inferences and construction at multiple (IP, subnet, and AS level) granularities. To eliminate redundant node comparisons, he uses a “just-enough-resolution” approach embedded in a tool called *rtd* (for route-diff), which compares (and quantifies differences in) paths at a finer (subnet or IP) granularity only if there are no differences at the AS-level granularity, i.e., to try to capture additional intra-AS topology information. With simulation and analysis of Ark+iPlane data, he demonstrated that his tool can provide powerful efficiency improvements, reducing computational and storage requirements which facilitates broader use, search, and sharing of the resulting data.

4. MOBILE MEASUREMENT

We spent the second morning listening to and discussing seven presentations on mobile measurement systems and tools. Yuanyuan Zhou (U. Michigan) talked about *Mobiperf*, an open source application for measuring network performance on Android mobile platforms (<http://www.mobiperf.com>). *Mobiperf* attempts to measure network throughput and latency, and supports other active measurements such as traceroute, ping, and DNS lookups. It uses Google’s M-Lab servers as test destinations as well as to share resulting data. *Mobiperf* outputs text format for users and json for researchers. As of February 2013 they had 600 active users, and 6GB of data. The big challenge is efficient data management and analysis. They no longer support the iPhone version because IOS does not allow background measurement. With the Android version the user can download the code, modify it to add new measurements, and reload it into the phone. Participants wondered how *Mobiperf* was getting accurate signal strength from what is usually noisy data. Aaron (UMD) noted that putting the phone in debug mode will yield a lot of information, including more accurate signal strength measurements. Yuanyuan said that they are working with Google but not directly with the Android team yet. They are interested in suggestions on what to do with the data collected.

Sachit Muckaden (Georgia Tech) talked about *MySpeedTest* (not to be confused with commercial product *SpeedTest*), a tool that actively probes five servers in the U.S. and Europe (Atlanta, Napoli, California) to measure TCP throughput, round trip time, jitter and loss. They plan to transition soon to using M-Lab servers as measurement targets to improve measurement accuracy. To minimize performance impact on the user, the tool only measures throughput on demand; it measures loss every two hours. The tool also passively collects network usage of data of applications on the device, and periodically (every 15 minutes) records other metadata that could affect user experience, such as signal strength, service provider, connection type, battery state, device type, manufacturer, time of day, and location. *MySpeedTest* is available on the Google Play Store and in February 2013 had 900+ active users from 115 countries who had run 1.5M measurements. Sachit has begun to correlate performance metrics with user behavior for different applications, looking for quality metrics that might help developers and service providers tune performance. They are working with M-Lab to try to share some subset of the data. Their biggest challenge is validating measurements against ground truth.

There was lively discussion after this talk on the goals of mobile measurement platforms. The original *SamKnows* broadband measurements (discussed at AIMS2012) were focused on verifying that users were receiving their contracted service. Another option is to have an application give warning, based on network conditions, about the impact that a given application will have on battery con-

sumption. John Heidemann (USC/ISI) posed a motivating question: in five years will wireless infrastructure performance have improved sufficiently such that sockets are still a good abstraction, or will performance be highly variable and force us to find a new abstraction? Aaron thought it depended on what regulators do with the spectrum. Ethan expressed concern about scaling these architectures beyond a few hundred users.

David Choffnes (U. Washington, soon to be Northeastern) introduced a new open source platform being launched by Google and U. Michigan (with Morley Mao at Google on sabbatical) to support participatory mobile measurement tool design and implementation. Like *SamKnows* for wired broadband, one goal is to gauge whether users are getting their expected levels of service, and even to identify a better-performing or lower-cost carrier for a given usage profile. Since validating deployed tools in the field is a daunting challenge, and running multiple apps doing the same type of measurements on a given device will reduce the accuracy of results, they are pursuing another approach: developing a open-source library of validated measurement primitives that everyone can use. Google’s appengine server provides data management, and sends an anonymized feed of data to M-lab.¹ Open problems include managing phone resources efficiently (battery life), and scaling to thousands of users. The library is implemented for the Android platform. David solicited feedback on measurement primitives to add to the library. Aaron noted that actively measuring a bandwidth-constrained system could itself cause a problem, which the tool should detect and warn about. Accurately identify this interference would require validation against ground truth from mobile carriers, which they have not yet tried. Nick noted that Comcast (at least) does traffic shaping above a certain threshold of congestion, but no one was aware that mobile carriers were doing traffic shaping yet.

Ethan (Yihua) Guo (U. Michigan) presented his group’s work on predicting mobile throughput based on device context such as signal strength and TCP connection state, which can impact battery life as well as user experience. They performed measurements to explore correlations between throughput and device context, with an aim toward using the data to more efficiently schedule network data transfers and reduce energy consumption on mobile devices. Consistent with the other performance measurement projects, the biggest challenge is validation. Unfortunately the tools they used gave answers with a wide range, from 8MB/s to 35MB/s, suggesting the tools are likely measuring different things. They are still considering how to share the data, and will need to anonymize the IMEI (unique identifier of phone) first. Aaron asked if signal strength determined maximum throughput, in which case a measurement could reveal that a user should not even try a certain application. Yihua believed the two were correlated but how tightly was not clear.

Ahmed Elmokashfi (Simula Research) gave an update on Norway’s mobile broadband measurement project presented at last year’s workshop. To address the lack of infrastructure to support mobile broadband measurements, Simula began a collaboration in 2010 to build a 90-host platform in ten municipalities across Norway to actively measure performance of three mobile broadband operators. They are now deploying more capable measurement nodes across a wider area (300 expected by March 2013), measuring five mobile networks, and storing data annotations including operator, geolocation, cellular identifier, and network modes. They create annual reports to help consumers choose operators. The infrastructure is partially funded by two mobile operators which may inhibit data

¹Measurement data available via a google account at <http://openmobiledata.appspot.com>.

sharing, but they have made the software running on the measurement nodes for managing multiple network interfaces open source (<https://github.com/kristrev/multi>) and they are considering making the platform accessible to researchers for experiments. This project may be the largest multi-carrier mobile measurement infrastructure. Aaron asked if the throughput drops during busy periods, but Ahmed said they do not measure throughput, only latency, which is definitely correlated with number of devices using the tower.

Džiugas Baltrūnas talked about the challenges in storing and representing mobile measurement results in a database. Performance measurements result in not only multiple performance parameters (delay, jitter) but also metadata valid only at an instant, e.g. signal strength, cell id, RRC state, network mode and submode, etc. Džiugas described how they use a relational database to efficiently store and organize this data. He described the structure of the tables, how to reuse tables for different measurements, and how standard database features can be used to extract metadata-enriched results. He solicited suggestions for experiments and how the data structure can be optimized further.

David Choffnes talked about Meddle, a collaboration with INRIA to increase transparency and control of mobile networks. Meddle relies on the fact that all major phones and carriers support VPNs, which allows tunneling to a server controlled by the project (not the carriers). Their setup allows for passive measurements on the phone, as well as device-wide ad-blocking. As of February 2013, they captured full packet traces for 19 devices across 14 users. Among the interesting preliminary results are that all Safari Google searches pre-iOS6 were in the clear, and that traffic from these users is split 60/40 wifi/cell, with little opportunity for compression. They hope to ramp up to 1000 users, using 10 tunneling servers in Seattle, Berkeley, Inria in France, and soon several in China. He invited those interested in who is tracking their cell phone web surfing to apply to participate in this IRB-approved study. But he also acknowledged that data sharing is not currently being considered given the sensitivity of the data, especially the full packet traces, although the next phase will not include full packet traces. Workshop participants suggested that they at least share the IRB application itself.

5. IPV6 ANNOTATIONS

We had three brief talks on IPv6-related topology issues. Billy Brinkmeyer (Naval Postgraduate School) spoke on how to infer router-level IPv6 topologies using a fingerprinting-based IP address alias resolution technique (PAM2013) that relies on inducing fragmented responses from IPv6 router interfaces. In IPv6, only devices originating packets typically fragment, not intermediate hops. They tricked a remote router into originating fragmented packets, by initially sending it an ICMP6 echo request, ignoring the reply, and then sending the router an ICMP packet-too-big message followed by new ICMP6 requests, forcing the router to fragment its replies. Using this trick they could measure how many live IPv6 interfaces responded, and how (e.g., with sequential fragment identifiers), to infer that two interfaces are likely on the same router. Their $O(n^2)$ algorithm demonstrated perfect inference accuracy in a controlled (emulated using real router images) environment, and on a small subset of the production IPv6 Internet for which they have ground truth. They performed Internet-wide testing and found that over 70% of IPv6 interfaces probed responded to the test. The biggest open challenge is getting additional ground truth to further validate the algorithm. In particular, what causes the other 30% to not respond is not understood; intermediate firewalls or routers may be choking on fragments. Verification would require more real-world testing with access to operators of the routers failing

to respond to TBT. So these results do not mean IPv6 alias resolution is solved, only that this technique constitutes a reasonable way forward, and a method to compare against other techniques. Since CAIDA will use this technique for alias resolution of its Ark IPv6 Topology Data Kit, they will share their results via CAIDA's data sharing mechanisms. Matthew briefly described a follow-on collaboration with Billy to scale his alias resolution technique to Internet-sized topologies.

Robert Beverly (NPS) talked about a different type of IPv6 inference: determining whether a given IPv4 and IPv6 address pair are assigned to the same interface, device, or even machine cluster. In collaboration with Arthur Berger (Akamai) and Nick Weaver, he described active and passive techniques aimed at associating addresses of network infrastructure, specifically DNS resolvers and web servers. Associating the IPv4 and IPv6 address(es) of dual-stacked DNS resolvers is important to Content Distribution Networks to leverage existing IPv4 reputation and geolocation databases for IPv6, while understanding the relationship between IPv4 and IPv6 addresses of web servers has implications on current cross-protocol performance measurement efforts. Further, given the relative lack of security of IPv6, this work seeks to show the extent to which IPv4 infrastructure depends on IPv6, and vice-versa. They found that 34% of inferred addresses are one-to-one, increasing to $\approx 50\%$ when aggregating IPv6 addresses into /64s. Yet they also discovered complex cases, with interconnected sets of address pairs that span continents and hundreds of autonomous systems; complexity attributable to large resolver clusters and distributed caches as confirmed using active probing. Future work involves performing similar inference on routers to understand topological inter-dependence.

We squeezed in two lightning talks at the end of the day. Casey Deccio (Sandia) gave a short description of his preliminary work characterizing IPv6 capabilities of hosts on IPv4 (SMTP) blacklists. He was seeking feedback on his methodology before he tried to deploy his technique on a larger, collaborative testbed. Greg Cole gave a beautiful demo of the Argus-based flow statistics reporter he uses for traffic measurement on his IRNC network infrastructure project Gloriat. He repeated a call he has made in his own community for help maintaining a database that maps IP addresses to organizations whose usage he needs to report. He monitors traffic from 12 million IP addresses at 30K institutions from about 2300 ASes seen each month. He has observed extremely little (although growing) IPv6 traffic. Some of the code supporting his system is third-party proprietary, but he is trying to release as much as he can open source.

6. CENSORSHIP AND OTHER OUTAGES

After some discussion about what we have learned from the workshop thus far, we focused the first part of Day 3 on topics related to the detection of censorship, filtering, and outages. Emile Aben (RIPE NCC) presented recent work analyzing the impact of Hurricane Sandy, as seen by RIPE's Atlas measurement infrastructure. During Sandy, Atlas had 2500 hardware probes in 104 countries, hosted in 1200 ASs, 400 v6 ASs. Each probe executed measurements to the thirteen root servers, which are anycast-addressed on hundreds of nodes around the world, so out-of-band measurements such as traceroute or a special DNS query are required to reveal which node is reached. Although this set of measurement targets is not comprehensive, they observed clear changes in delay and paths from the effects of Sandy. Most of the path seen in this traceroute-based study occurred within an AS, so would be invisible at the BGP-level. The predominant path change observed was not in New York but toward the DC/Ashburn area. Surprisingly, two At-

las probes in NYC stayed up throughout the storm, while electrical power to most of the city was out for several days. (Data centers have generators, homes mostly do not.) This analysis convinced Emile of the importance of more accurate router geolocation; he used airport codes and city names in DNS labels where he could, but not all router names have usable geographic hints. The group discussed how to re-architect probing to capture these kinds of outages, while keeping the geolocation challenges and measurement load tractable. Emile emphasized the value of probing known locations that do not move, like Ark and PlanetLab nodes. Aaron suggested following news feeds regarding weather or other likely indicators of imminent outages, and reactively point probes at targets in those regions. Robert pointed out the possibility of measurements making a bad situation worse during times of reduced connectivity, so there is a trade-off. Nick suggested they use the 2000+ Ark nodes themselves which are at fixed locations, but they tend to be behind NATs which interfere with probing. Emile later presented RIPE's reporting of counts of v4 prefixes and ASs observable via RIPE RIS data, which shows the recent Syrian Internet outage: <https://stat.ripe.net/widget/country-routing-stats>.

John Heidemann (USC/ISI) talked about his long-term (since 2006) analysis of outages, using active pinging of edge networks. He discussed challenges in collection and analysis of long-term, general purpose datasets, and using them to improve outage detection via active probing. The biggest challenge is selecting networks to probe. Applying statistical population sampling, they used their IP censuses to walk the entire Internet address space periodically, and more focused surveys to probe more frequently and thoroughly a fraction of that space. Since a failure to respond from a single IP address in a /24 is an ambiguous result, the surveys selectively probe some /24s completely, which allows correlation of entire block failures to outages. In the case of Sandy, their probing completed one measurement cycle during the storm. Geolocating that data revealed that the number of outages in the NY/NJ area doubled relative to the baseline, and took about 4 days to recover. Historical data also revealed the effects of the earthquake in Japan, the Egyptian censorship episode, an Australia outage in January 2011, as well as other smaller events. Having a consistent longitudinal data set that establishes a baseline makes it easy to identify prominent changes. He uses a wiki to share data and results, and also uses the PREDICT project to share the data with other researchers. In the course of his probing he has dealt with complaints by maintaining a don't-probe-me-list and sharing that with other researchers too. They have not yet explored using this data for predictive models. Sometimes they adjust their survey start times based on expected or ongoing events, but they do not change the polling rates. Since the same machines are performing both the surveys and censuses, they have to be careful not to tweak parameters that may interfere with continuing measurement.

John-Paul Verkamp (Indiana U.) described his ideas to use DNS, BGP, and other data to infer censorship. They queried open resolvers in different countries, validating locally against non-censored results. His initial data sample suggests 73 of 202 observed countries are engaged in censorship or filtering, between a third and a half of those using DNS-based methods. Censorship implementations vary, from directly controlling DNS resolvers, cache poisoning, or in-flight packet modification or injection. For example, China's great firewall acts as a resolver for Facebook and Twitter. He does not have data to share yet, but he does have a list (incomplete) of keywords being filtered in China, and a list (also incomplete) of lying resolvers.

Phillipa Gill (Citizen Lab/Stony Brook) questioned Jean-Paul's results since she observed only 6 countries (where they had at least

50 measurements) doing DNS-based censorship more than 10% of the time. She also noted several challenges in characterizing global web censorship: gaining access to realistic vantage points in the country, informed consent from participants, and managing longitudinal data. Using data collected by the OpenNet Initiative (<http://opennet.net>), which has a rich longitudinal data set (2007-2012) from vantage points in 77 countries and 286 different ISPs, although it was not collected by network measurement experts. The basic idea is to issue requests for a consistent set of URLs (by informed and consenting users) from a field location as well as a control location (the lab), and compare results between the two (including DNS, HTTP headers, etc.) to identify instances of blocking. The project ran for six years before funding was reduced. The top six countries in terms of observed blocking were: China, Iran, UAE, Yemem, Burma, and Vietnam. She also saw variation across ISPs in the same country, across different types of networks (academic networks blocked less), and across different types of censored content (URLs with different types of content are blocked differently). She saw shifts in blocking over time, e.g., after political reform. She iterated several needs: more rigorous experimental design; a sustainable model for maintaining infrastructure; techniques that distinguish censorship from other outages. She hoped to index her data in DatCat when it is available. Mehmet asked whether the lab experienced hacking; Phillipa had heard about hacks against the lab but not against the volunteers doing the measurements. Her data did suggest that if China's great firewall is overloaded, it stops filtering.

Ethan Katz-Bassett (USC) gave a short talk on the BGP-Mux (aka Transit Portal) project, started by Nick Feamster and now jointly run with Ethan. This testbed is designed to handle the operational details of supporting experiments where the researcher wants to act as an AS, i.e., speak BGP and send/receive traffic. With nodes in five universities acting as providers, the testbed supports virtual links to experiments at these or other universities. The researcher looks like their own AS, with an AS number, using experimental IPv4 space. This platform allows researchers to measure and experiment with interdomain routing. Ethan used it for his *LIFEGUARD* project² (SIGCOMM2012) to locate Internet failures and dynamically generate usable alternate routes. Ethan and collaborators are also using it to analyze root causes of BGP path changes (SIGCOMM2013), and Nick is using it for his PECAN project, measuring performance of alternate paths in joint content and network routing scenarios (SIGMETRICS2013). Ethan has published his reverse traceroute data online via Google's cloud storage. An online map shows the testbed status (<http://tp.gtnoise.net>). Ethan and Nick are interested in expanding the testbed footprint and user base. When asked about whether he gets nervous while breaking routes, Ethan noted that they have not had any problems thus far due to careful filtering, such that routing changes only affect traffic to the experimental prefix, but they do notify NANOG when they're about to do something unusual.

Ramakrishnan Durairajan (U.Wisconsin) talked about his lab's Internet Atlas project, a new database and visualization portal of the physical interconnection structure of the Internet. They want to create a comprehensive and geographically accurate catalogue and analysis environment for: (i) the locations of points of presence (PoPs) that house switching and routing equipment, (ii) the conduits (links) that connect these locations, and (iii) relevant metadata, e.g., source provenance. They use web search and a library of parsing tools to capture maps and other provider topology information, extended with related data from active probes, BGP updates,

²*LIFEGUARD* stands for Locating Internet Failures Effectively and Generating Usable Alternate Routes Dynamically

twitter, weather, which they enter into a database using both manual and automated techniques to conserve consistency. They use Google's geocoder service to geolocate PoP addresses. The repository currently contains over 8K PoP locations and nearly 13K links for over 185 networks (including all tier 1 providers) around the world. They also have meta-data on 744 NTP servers, 221 traceroute servers, 358 IXPs, and DNS root name server instances. The openly available web portal is based on the widely-used ArcGIS geographic information system, which enables visualization and diverse spatial analyses of the data. They expect it can be used to look for vulnerabilities in the current infrastructure, as well as to seek opportunities to discover or improve peering arrangements. Young asked how they verify their data; Ram said they currently manually compare images extracted from ISP maps on web sites to the parsed form of the ISP map in the Atlas database.

Aaron Schulman (U. Maryland) reported follow-up work from last year, focused on pinging U.S. home users before and after a weather event. He showed an animation of his data from pinging 60K hosts in the U.S. every 11 minutes during Hurricane Sandy, which revealed correlated outages in other areas, e.g. in MD before Sandy hit the coast of New Jersey. He solicited (and received) speculation on how to interpret some discontinuities and gaps in his data (potentially effects of DHCP, rate-limiting, other carrier-specific behavior). He compared durations of problems across three providers: Verizon FIOS, Verizon DSL, and Comcast, and found that Verizon FIOS outages last longer, possibly because there is less experience with repairing this newer technology. He can sometimes distinguish power outages from network outages, but he noted the available data on power distribution is extremely messy. (After AIMS, Aaron's group discovered that the US Department of Energy publishes a list of known power outages³, so they were able to discover that a power outage appears as a correlated failure of at least two IPs in at least four ISPs.) He acknowledged that his inferences were only as good as MaxMind's geolocation accuracy, and he would love to have street-level geographical data.

Nick Feamster (Georgia Tech) spoke on his recent research that exposes inconsistent web search results with browser plug-in Bobble (<http://bobble.gtisc.gatech.edu/>). Although most search engines customize (personalize) search results based on a user profile, the algorithms they use to do so are opaque to users. Bobble is a tool that executes a single user query from different vantage points and parameters (using PlanetLab nodes), and compares results returned. Using more than 75K search queries from about 175 users over nine months, they found that 98% of all Google search queries from Bobble users resulted in some inconsistency, and that geography was more important than search history in affecting search results. Inconsistencies are pervasive – about half of the queries tested returned at least five unique result sets. He cannot share the results due to IRB requirements. Nick Weaver suggested that rather than using PlanetLab to run the same query from the same machine, to use instead a cookie-free browser that is not signed into any Google services. In the future they would like to find ways for users to correct bias in search results. Robert asked for a version of Bobble that does not intercept all queries.

Eric Osterweil (Verisign) gave a plea from industry to have more measurement researchers become aware of current operational challenges in interdomain routing security, specifically the framework being established by the IETF for Internet resource (e.g., IP addresses) certification. The Resource Public Key Infrastructure (RPKI) is a new standard that defines a distributed database to enable routed resource certification for secure inter-domain routing. There is a

vision for a single root trust anchor (TA) that will be managed by IANA to ensure no allocation conflicts. However, this goal has proven elusive, and we currently have 5 root TAs (one for each RIR), which allows allocation conflicts. The design of the RPKI envisions routed resource data (IP prefixes and ASNs) shared between authoritative repositories and client side caches. About 1% of today's BGP-routed prefixes are registered in an RPKI repository, at least one of which has itself already suffered a few major outages. The entire RPKI is a cryptographic delegation chain with eventually potentially hundreds of thousands of objects, but since the RPKI is intended to inform BGP routing processes in near real-time, routers would need keys and to sign/verify updates. There is concern that the RPKI will not be able to keep up with the speed requirements of routers to process RPKI updates. Today a routing change takes minutes; Eric fears that with RPKI+BGPSEC it could take days (which would interfere with several companies including Verisign who do re-routing to protect customers being DOS attacked). The transfer protocol in the standard uses rsync, which itself has scaling challenges. Other issues include: RIRs could override each other, including potentially inducing surgical takedown of another RIR's resources; projects like BGP-Mux become infeasible; and the current RPKI+BGPSEC design does not address the problem of route leaks. Eric has co-authored a report outlining scaling challenges faced by the RPKI (published at <http://techreports.verisignlabs.com>). His lab also supports a tool that tracks updates to the RPKI and tweets prefix/origin binding changes. It is not yet clear what happens to the legacy IP addresses in this model, since they have no authoritative RIR. But FCC advisory groups are currently trying to decide what to recommend for resource certification, so it is critical to raise awareness now, especially by empirically-minded researchers.

To end the technical session, Young Hyun (UCSD/CAIDA) reported on two big updates to CAIDA's Ark active measurement infrastructure: support for on-demand measurement, and deployment of Raspberry Pi hardware as new Ark nodes. The Ark monitors all ping and traceroute operationally, and CAIDA now has a (non-public but accounts available to researchers upon request) web interface to request measurements of the Ark monitors. One can request, e.g., for all Ark monitors to measure RTT and traceroute to a single destination. The second update is a gradual transition from 1U hardware to support Ark nodes to a cigarette-pack-sized \$35 Raspberry Pi node. CAIDA has deployed 5 as of February 2013, mostly outside the U.S.

7. DATA SHARING

Erin Kenneally (CAIDA/UCSD) led a discussion on data sharing issues culled from workshop conversations, especially the risks related to collection and disclosure of mobile measurement data. She emphasized the importance of clearly articulating the need for and benefits of sharing, e.g., to parameterize or validate models and inferences, to avoid the cost of duplicate collection, and to establish historical baselines of network behavior. The cost and risk of sharing data is the starting and often stopping point for decision-makers, which means their default-deny positions must be actively counterbalanced. Their risk aversion is reinforced by highly-publicized cases (e.g., AOL, Netflix) where anonymized data has been de-anonymized in embarrassing ways, as well by the fact that in general the application of privacy is immature in the network data realm.

Participants discussed examples of complications giving rise to real and perceived risks. RIPE is sharing the Atlas data, but Robert worries about risks because probes are in different countries with different rules and expectations. Nick gave an example of the sen-

³<http://www.oe.netl.doe.gov/oe417.aspx>

sitive information that can be extracted from Netalyzer data, e.g., what ports are blocked by military hosts, and emphasized the difficulty of assessing the risk of releasing seemingly benign network data when it can be combined with other readily available data. Nick described how ad hoc data sharing arrangements using a code-to-data model can emerge as sharing parties better communicate and formalize what they want, e.g., a DB schema that gets vetted for sensitivities before data is released. Nick acknowledged that putting Netalyzer metadata in DatCat or PREDICT would help let researchers know about the existence of this (send-DB-scheme-to-data) sharing technique.

Nick also highlighted ISC's SIE project as a successful example of providers sharing data. He also noted that some researchers are "over-sharing data", highlighting the case of Indiana University making a huge volume of edge web traffic data readily available, thereby greatly enhancing the chances that harmful uses could arise. Erin asked if researchers or data providers saw a benefit to a third party (proxy) that sanitized or otherwise mediated the risk-sensitive disclosure the data, but Nick thought such a role would not offload the risk. Erin agreed but noted that it could alleviate cost concerns, and move participants closer to a shared risk model. Erin mentioned the Menlo Report released last year, a proposed set of guidelines for balancing risk and benefit in Internet research. She is chairing a related IEEE workshop, "CREDS: Cyber-security Research Ethics Dialogue and Strategy" (May 2013), which will include discussion of research using data that is publicly available but of questionable lineage.

The strongest potential data sharing collaboration that emerged from the discussions was a crowd-sourced geolocation database that could improve on the accuracy of infrastructure geolocation, i.e., routers and PoPs that the commercial services are not trying (hard) to geolocate accurately.

8. RESULTING COLLABORATIONS

Several collaborations were continued at the workshop, including those that originated at previous AIMS workshop. New collaborations initiated included:

1. Google/M-Lab (Dominic), Georgia Tech (Nick + Sachit) and David Choffnes have started discussions to merge their mobile measurement library efforts.
2. Dave, Morley, Matt Welsh, Ramesh, students, and Ethan are collaborating on mobile measurements.
3. David Choffnes briefly discussed with Georgia Tech about combining measurements from Meddle and Bismark.
4. John Heidemann and Dan Massey plan a collaboration with Greg Cole on Gloriad network traffic data analysis.
5. Ethan's student will travel to Amsterdam to deploy a new BGP-Mux node at AMS-IX to peer with 500 ASes there, and hopefully stop by RIPE to discuss collaboration.
6. Rob Beverly and Casey Deccio are extending Casey's analysis of IPv6 capabilities of suspicious IPv4 hosts. Rob visited Sandia after the workshop to discuss his transport-layer traffic analysis work, which Sandia may use internally.
7. RIPE (Robert, Emile) and CAIDA are interested in pursuing a crowd-sourced geolocation database, and possibly holding a workshop later this year focused on this topic.
8. RIPE (Emile) is looking into whether Sandor Laki can do triangulation measurements on RIPE Atlas to support his spatial analysis of topology.

ACKNOWLEDGMENTS. The workshop was supported by the National Science Foundation (CNS-0958547, OCI-1127500, and OCI-

0963073) and by the U.S. Department of Homeland Security (DHS) Science and Technology (S&T) Directorate (N66001-12-C-0130). We thank all participants for their insights, feedback, and contributions of text to the report, and many thanks to Evi Nemeth for taking excellent notes at the meeting.