

The 11th Workshop on Active Internet Measurements (AIMS-11) Workshop Report

kc claffy
UCSD/CAIDA
kc@caida.org

David Clark
MIT/CSAIL
ddc@csail.mit.edu

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 It has NOT been peer reviewed.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

ABSTRACT

On 16-17 April 2018, CAIDA hosted its eleventh Workshop on Active Internet Measurements (AIMS-11). This workshop series provides a forum for 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 data. An overarching theme this year was scaling the storage, indexing, annotation, and usage of Internet measurements. We discussed tradeoffs in use of commercial cloud services to make measurement results more accessible and informative to researchers in various disciplines. Other agenda topics included status updates on recent measurement infrastructures and community feedback; measurement of poorly configured infrastructure; and recent successes and approaches to evolving challenges in geolocation, topology, route hijacking, and performance measurement. We review highlights of discussions of the talks. This report does not cover each topic discussed; for more details examine workshop presentations linked from the workshop web page:

<http://www.caida.org/workshops/aims/1904/>.

CCS CONCEPTS

• **Networks** → **Network measurement; Public Internet; Network dynamics;**

KEYWORDS

active Internet measurement

1 STORAGE AND USE OF INTERNET MEASUREMENTS

The field of network measurement continues to yield innovative and creative techniques and methodologies, but has also expanded in scope to related challenges of efficient storage of large quantities of measurement results, and improving accessibility of measurements to facilitate analysis. These two challenges are coupled: the approach used to store the data will shape or limit the types of queries that can be implemented with reasonable performance and without specialist training in the query tools. Several projects are exploring commercial cloud service platforms as resources to store data for community use. We discussed the trade-offs inherent in such choices.

RIPE Atlas. RIPE NCC has over 10K Atlas probes in operation. They have over 400 “anchors”—more powerful devices that serve

as destinations for some experiments and the basis for more complex measurements. The RIPE NCC currently hold about 66TB of compressed measurement data, and operates an internal, private, Hadoop cluster for analysis of accumulated measurements and metadata. Query execution can be slow, and maintenance is expensive, relying on multiple specialists to manage software and hardware systems. For external users, public results are available in bulk via FTP, or via an API, but with no query interface. Stephen Strowes (RIPE NCC, visiting CAIDA this year) reported on a new experiment led by Elena Dominguez (RIPE NCC) to put RIPE Atlas measurements in the Google Cloud (BigQuery) Platform, to improve query performance, improve service availability, and reduce operating costs, for both internal (to RIPE) and community use. CAIDA is pleased to be hosting Stephen Strowes this year to exchange experiences with storage, curation, and indexing large volumes of archived traceroute data for research purposes. He will also undertake a quantitative comparison of the diversity in the IPv4 and IPv6 topology data collected from each (RIPE Atlas and CAIDA's Archipelago) platform, in terms of vantage points, targets, and intermediate hops, along the following granularities: BGP-routed prefixes, networks (AS-level), network types, and geographic coverage.

Route Views. David Teach (University of Oregon / Route Views), described their plans to modernize and upgrade the Route Views platform. The Route Views project has been in service for 24 years, the longest-running of any global Internet measurement infrastructure, and holds about 22 TB of compressed information reflecting 833 peering systems across 239 unique ASes peering with Route Views. The current pipeline is based on open source software, including FRR (a fork of Quagga with more commercial support), OpenBMP, and Multi-threaded Routing Toolkit (MRT). Route Views hopes to evolve the collector infrastructure to support real-time data delivery, and new metadata such as RPKI usage. Route Views is also interested in governance models to be able to adapt to the future needs of the research community. One request from the community was the ability to allow peers to export their entire RIB (not just FIB) to Route Views via BMP.

Censys. Zakir Durumeric (Stanford University) discussed the history and evolution of his team's Internet scanning work. Over the past few years, they have released several tools, datasets, and services to share scanning data with security researchers: the ZMap scanners, Scans.IO Data Repository, and Censys Search Engine. This project started as small academic research effort in 2012, but

in a few years outgrew its university infrastructure capabilities. In 2017 Zakir started a for-profit company, Censys, which by 2018 was serving over 1.5 million user queries per day from 90K registered users, accumulating 1-2 TB of new structured data per day. Zakir reviewed many lessons learned from their experience. They are moving away from the approach of designing web interfaces or APIs, since they have concluded that a better approach is having users query the data via BigQuery with Google's cloud infrastructure as a backend. They also found that re-writing Python modules into Go was easier than finding memory leaks in other people's Python modules! They found that JSON's lack of format constraints complicated data processing; they chose to use Protobuf and Avro instead. They found Elasticsearch was expensive to scale (they deployed 48 hosts for 20 TB) and struggled if given a poorly crafted query (perhaps with a costly regex). Many in the room hoped that Zakir's team would write a write paper expanding on these lessons learned, since many other academic measurement infrastructure projects encounter similar challenges.

Commercial flow data analysis service. Avi Freedman (Kentik) talked about how Kentik, a commercial software-as-a-service platform, enriches traffic flow measurements with additional business, routing, application, and security context, enabling operators to drastically increase their ability to reason about traffic, performance, and attack patterns. He presented a brief overview of Kentik's scalable architecture for enabling real-time syncing with metadata, DNS, and routing systems, performing streaming joins to enrich traffic data, and sending the serialized output to multiple systems for storage, querying, and anomaly detection. He was interested in opportunities to work with researchers who can extract useful insights from aggregated data that his customers are willing to share.

Platform for Applied Network Analysis (PANDA). To support complex queries on multi-dimensional data, the commercial cloud offerings are not (perhaps yet) providing sufficient tools and platforms to obviate the need for custom platforms. CAIDA researchers talked about several such platforms, and components of future platforms. kc claffy (CAIDA/UC San Diego) reported on CAIDA's effort to design and build an infrastructure that will allow scientists, operators, and eventually policy makers to ask high-level cross disciplinary questions about Internet structure, behavior, and performance. This NSF-funded project¹ is integrating several data building blocks that CAIDA researchers have developed over the last twenty years: the Archipelago (Ark) Active Internet measurement platform and its derivative data and analytics components; ASRank, which computes and compares routing and economic relationships among ISPs; BGPStream, an efficient framework for routing (BGP) data analysis; MANIC: Mapping and Analysis of Interdomain Congestion (described below); Periscope, a unified interface to looking glass topology measurement servers; and Spoofer, a system for assessment of IP source address validation (a best practice) compliance. The goal is to develop unified mechanisms to support access to data sets generated by these platforms, to facilitate collaboration with and across multiple disciplines, and increase

¹NSF's Data Infrastructure Building Blocks program: Integrated Platform for Applied Network Data Analysis (PANDA), <https://www.caida.org/funding/dibbs>.

community accessibility to the underlying components and data. The project targets research questions from four disciplines: networking, security, economics and policy.

We reviewed the state of development of two PANDA components: Archipelago (Ark) Internet measurement platform APIs, and MANIC. CAIDA also presented updates on their *IP address alias resolution* services – the process of identifying which IP interfaces belong to the same router, in order to infer a router-level topology. Young Hyun (CAIDA/UCSD) described recent developments for *aliasq*, a web API for querying a database of aliases harvested from CAIDA's Internet Topology Data Kits.

Alex Marder (UPenn) discussed how to use traceroute data to more accurately infer layer-3 topology with virtual private network (VPN) technology (in particular Virtual Router Forwarding, or VRF) layered on top. A VRF is a virtual routing table that corresponds to a VPN and participates in BGP route announcements, but is separate from any other routing table on its router. Failure to account for VRFs in traceroute can lead to mistakes in Internet topology analysis, such as inferring router ownership and identifying interdomain links. Alex introduced preliminary work on detecting VRF forwarding addresses in traceroute, using active and passive approaches, and showed that these techniques perform accurately compared to ground truth from two research and education networks.

MANIC: Measurement and analysis of interdomain congestion. Roderick Fanou (CAIDA) described CAIDA's Measurement and Analysis of Interdomain Congestion (MANIC) system (<https://manic.caida.org>), which CAIDA prototyped to monitor interdomain links and their congestion state, in order to provide empirical grounding to debates related to interdomain congestion. The project focused on interdomain links because they represent points of interaction between ASs, where business issues may lead to poor traffic management and capacity planning. MANIC is currently deployed at 86 vantage points worldwide and has collected data since 2016. CAIDA has shown that congestion inferred using the lightweight TSLP method correlates with other metrics of interconnection performance impairment. CAIDA researchers and collaborators used this method to study interdomain links of eight large U.S. broadband access providers from March 2016 to December 2017, validating their inferences against ground-truth traffic statistics from two of the providers. CAIDA has built and released access to the MANIC API (<https://manic.caida.org/>). Roderick Fanou (CAIDA) presented the state of the API, its architecture, and demonstrated its functioning.

UCSD Network Telescope Data. Alistair King (CAIDA) talked about CAIDA's new STARDUST platform² to modernize the UCSD Network Telescope's capture infrastructure (from 1GB to 10GB), including a dedicated DAG for accurate timestamping, capturing packets, and immediately spitting them out to a multicast VLAN to connect researchers to a realtime traffic stream. The STARDUST project will scale storage and access using virtualized and containerized environments, allowing each user to start in a smaller container and scale up compute needs nearly transparently to the researcher. A goal is reduction of the "time to insight", by virtue of the lower bar for researchers to access data. The platform will support processing

²<https://www.caida.org/funding/stardust/>

of millions per packets per second, 100s of GBs / hour, creation of high-level aggregated data sets, and richer annotations.

2 GEOLOCATION

A recurrent theme of discussion was geolocation of Internet elements, e.g., IP addresses, domain names, routers, links.

Geolocation of IP addresses. While current commercial services that map IP address to location work reasonably well for end-node addresses, they are much less reliable for addresses associated with routers toward the center of the network. Robert Kisteleki (RIPE NCC) spoke about RIPE’s IPmap geolocation project,³ which is exploring a new technique to geolocate an IP address, which they call *single radius*. To execute this technique, they use a subset of their probes to measure the latency to a target address, and iterate with presumably closer probes until they find a probe with a sufficiently low RTT that it must be near the location. The system supplements this *single radius* scheme with crowdsourced location assertions.⁴

Country-Level IP Geolocation. Ioana Livadariu (Simula Metropolitan) compared recent methods and databases that map IP addresses to countries. They found that RIR delegation files and commercial services like MaxMind and IP2Location had high coverage of IP addresses along traceroute-observed paths, whereas active-based geolocation datasets like RIPE IPmap’s single-radius technique covered at most half of such IP addresses. They found The three databases disagreed in geolocating about 10% of the measured IPs they examined, and they traced most disagreements to mergers and acquisitions between ISPs. IPs that belong to geographically disperse networks were more likely to geolocate differently across databases. While geo-hints encoded in the domain names could mitigate these limitations, they found many cases where the DNS-based geo-hints communicate misleading information.

Cartography of physical elements. Justin Rohrer (Naval Postgraduate School) talked about *net.tagger*, a crowdsourced approach to mapping physical communication infrastructure. Currently, data on city, metro, and street-level deployment of Internet infrastructure is scarce, making it hard even for providers to understand how they share fate due to use of shared fiber conduits, towers, etc. The *net.tagger* project intends to fill these gaps by crowdsourcing the tagging of local telecommunications infrastructure via a simple-to-use smartphone app. Users photograph evidence of physical infrastructure and upload the data to an OpenStreetMap-compatible database. Analysis techniques include validation and connecting the dots between tags through pattern inference.

Inferring Country-Level Transit Influence of Autonomous Systems. PhD student Alex Gamero-Garrido (CAIDA/UC San Diego) discussed the challenge of identifying the most influential transit providers in each country, i.e., those ASes that may have the potential to observe, manipulate or disrupt Internet traffic flowing toward that country. CAIDA has developed two new Internet cartography metrics to overcome several challenges with making such inferences using BGP data. The transit influence (TI) metric for a

transit AS estimates the share of addresses of origin AS served by the transit AS. The Aggregate Transit Influence (ATI), captures the aggregate of all fractions of each country’s origin ASes’ addresses that the transit AS serves. Alex described how to apply these two metrics to identify the most influential ASes in each country, and the origin ASes in those countries that heavily depend on transit ASes.

3 SECURITY

BGP poisoning. BGP poisoning is the act of announcing an AS path to a prefix that artificially includes one’s own AS number. For example, to poison the incoming route through AS 2, AS 1 might announce a path that includes the AS sequence 1,2,1,... AS 2 would reject this path because it appears to contain a loop. The owner of a prefix can use this poisoning technique to avoid receiving traffic from AS 2. Jared Smith (U. Tennessee /ORNL) described a study to demonstrate the feasibility of BGP poisoning as a security primitive, with applications such as censorship circumvention or defense against DDoS. He created an overlay testbed measurement system spanning 5 BGP routers, 8 previously unused IP prefixes, and 5,000 vantage points across 3 countries. They also explore the prevalence of filtering that prevents poisoned path propagation. In their measurements, over 80% of observed ASes from route collectors propagated poisoned paths.

BGP Hijacking Observatory and Mitigation. Alberto Dainotti (CAIDA/UC San Diego) introduced CAIDA’s new BGP Observatory project, which monitors the state of BGP routing using RIPE and Route Views data to detect anomalies that may be episodes of BGP hijacking. For each suspicious event, the system augments control-plane (BGP) data with data-plane measurements (traceroutes) executed from a set of RIPE Atlas probes. Events are enriched with descriptive tags based on various database lookups (e.g., AS relationships, AS Customer Cone, IXP prefixes) and heuristics (e.g., potential “fat finger” misconfiguration) and presented to users through a Web interface that facilitates inspection of events. The observatory serves multiple purposes: a platform for operators to troubleshoot anomalous events and enable situational awareness, and a testbed to realistically experiment with detection techniques applied to Internet routing data in the wild. Alberto also described a related recent collaboration to develop a tool (ARTEMIS) for an enterprise to detect BGP hijacks of its own prefixes. A network owner installs ARTEMIS to observe legitimate and anomalous announcements about their prefixes observed from other parts of the Internet. This tool depends on an infrastructure for gathering and processing route announcements, and allows detection of a route hijack within seconds. Since the tool has access to the ground truth about legitimate assertions, it can automatically trigger remediation, including fragmenting the legitimate announcement into more specific prefixes that override the malicious announcements. Currently ARTEMIS is under pilot testing at several network operators, including Internet2.

Measuring poorly configured infrastructure. Matthew Luckie (U. of Waikato) talked about two crowdsourced measurement projects

³<https://ipmap.ripe.net/>

⁴https://labs.ripe.net/Members/massimo_candela/ripe-ipmap-whats-under-the-hood

⁵<https://export.arxiv.org/abs/1811.03716>

to capture security hygiene properties of networks on the global Internet. The first project, Spoofer, is a client-server system that tracks compliance with source address validation practices⁶ to prevent networks from being used to launch spoofed-source-address denial of service attacks. This project’s focus has shifted toward understanding the impact of various approaches to remediation, e.g., notifying administrators of networks that fail SAV compliance tests, and automatically sending reports to regional public network operator mailing lists of all networks in the region that failed a compliance test, and those that remediated, in the the past month. The second (and much newer) project, Netstinky,⁷ provides an iOS/Android app client that notifies the user when they are attached to a network listed on one of the public blacklists for exhibiting evidence of malicious conduct.

Gregory Petropolous (Security Scorecard), described how his company measures corporate security at scale. To assess the “security health” of a desktop, they purchase user agent strings from ad network data, and map these strings to their origin, allowing them to assess over 1 million companies with respect to metrics such as how quickly software updates are installed, or whether they are running versions of software with known vulnerabilities. They use this data as a component of an overall security score. This data is sold to companies all over the world to monitor their own cyber risk as well as the risk that their vendors introduce.

Evading Nation-State Censorship with Genetic Algorithms. Kevin Bock and George Hughey (U. Maryland) described how they used genetic algorithms to automate the process of circumventing censorship. Today’s evade/detect cycle is largely manual: researchers actively measure censoring networks to learn how they operate, and then develop strategies to exploit shortcomings in the censors’ designs and implementations. This cycle gives censors an advantage: it takes longer to learn about censors than it takes a censor to adapt its methods. Bock and Hughey’s novel approach uses artificial intelligence to adaptively probe censoring regimes and automatically discover strategies for circumventing them. The general technique used by their circumvention tool is to send packets that disrupt or confuse the state stored in the censor’s monitoring box. As a simple example, if the sender transmits a TCP Reset with a TTL that causes the monitoring box but not the receiver to see it, the monitor will conclude that the session has been reset but in reality the session can continue. Automated genetic-based exploration of the algorithms used by censorship devices re-discovered all the prior known methods, and found 4 unique species of HTTP strategies that could evade Chinese censorship.

4 PERFORMANCE MEASUREMENT STUDIES

Several other workshop presentations touched on issues of macroscopic performance measurements.

Why Anycast Routes Aren’t Good. Neil Spring (U. Maryland) described experiments to understand anycast routing decisions across ASs in the Internet. The University of Maryland operates one of the root DNS servers, which has many replicas addressable using an anycast address. This system allows them to explore how queries to

their root service from various sources travel to a particular anycast instance of their service. They found that in practice, anycast is particularly vulnerable to bad BGP routing decisions. Distant replica are often selected, for reasons that are not easy to discern. There is a more general lesson here about the challenges that interconnected ISPs face if they want to offer more advanced services such as “routing to a service”. There will be a need for a higher degree of cooperation and information sharing, which is challenging for competing ASs to implement.

A First Look at the FCC’s Measuring Mobile Broadband Data. There are few publicly available studies of mobile broadband performance based on large-scale datasets. Vendors of some mobile speed test apps occasionally release reports based on their data, but the measurement methodology is highly variable and often proprietary. Typically such reports lack detail on distributions of performance metrics, packet loss, or performance variance over time. Scott Jordan (UC Irvine) presented initial analysis that he and James Miller (Amazon) have performed on the FCC mobile speed measurement data. The FCC has been collecting data since 2013 on the performance of mobile broadband Internet access service through its Measuring Broadband America program, but only recently published a set of raw data (with no accompanying report or analysis), the first large-scale publicly available dataset on U.S. mobile broadband performance. The raw data provides a view of performance across major providers based on a uniform measurement methodology. However, there are many questions about how this data should be analyzed and presented, and perhaps a bigger question of who would have incentive to objectively do so. Not surprisingly, there is high variance among different measurements of the same provider, so a mean or median is not a sufficient representation of the data. Preliminary analysis of the variance in speeds from different providers revealed significant disparity between what is observed and what is advertised.

Evaluating accuracy of web speed test platforms. Web-based speed test measurement platforms use HTTP or WebSockets to flood the network as a way to measure available bandwidth in both downlink and uplink directions. The rapid increase in residential broadband access link capacities over the years presents new challenges for measuring available bandwidth. Ricky Mok (CAIDA/UC San Diego) is leading a project to appraise the accuracy of five major speed test platforms — Ookla Speedtest, M-Lab NDT, Comcast Xfinity, Netflix fast.com, and speedof.me. In addition to conducting speed tests against these platforms from clients in different networks and locations, he is using information from packet traces and the browser to conduct cross-layer analysis. His preliminary finding is that the maximum capacity that these popular speed test platforms can accurately measure is generally far lower than 1Gbps.

Crowdsourcing QoE measurements. Ricky Mok (CAIDA) and Ginga Kawakuti (NTT) gave a status update on their collaboration to build a platform and framework to crowd-source measurements for assessing Quality of Experience (QoE).⁸ With support from NTT, the Quality of User Internet Customer Experience (QUINCE) project will examine the feasibility of gamifying various kinds of Internet

⁶<https://spoofer.caida.org/>

⁷<https://netstinky.wand.net.nz/>

⁸<https://crowdtrace.caida.org/>

measurement and subjective assessments of video streaming performance by using a web-based platform. The project will test the sustainability of performing these measurement tasks in a crowd-sourcing context and will collect data to study the topology and performance of the Internet, the streaming performance of video services, and subjective assessments of the quality of experience (QoE) of video streaming.

Passive measurements of performance. As access speeds increase, it is less likely that the access link will be the throughput bottleneck, and thus active speed measurements will not be good predictors of application quality. Additionally, active probing that by design fills up the access pipe becomes more disruptive as access speeds increase. Renata Teixeira (Inria / Stanford) described her collaboration to use passive measurement to study residential Internet performance. Passive data gathering generates huge data sets, so a critical part of this research may involve sophisticated approaches to real-time aggregation and processing, and of course managing privacy issues. This collaboration has developed a passive measurement infrastructure that is deployed in dozens of homes to capture packet traces. The device will be used to infer the quality of video playback, based only on access to the encrypted data stream, together with captured DNS queries that allow them to map flows to specific content providers. They have access to ground truth data from Netflix, Youtube, Twitch, and Amazon Prime, which they have used to train their inference engine. One preliminary result is that the capacity they infer from their analysis is a better predictor of video quality than the advertised access speed. They plan to investigate the capabilities and limitations of passive monitoring to measure application performance and explore how targeted active measurements can help locate the source of application performance bottlenecks.

Passive measurement to detect congestion. Large content providers have network management systems that can track congestion within the region of the network that they control, including points of interconnection into adjacent networks. However, if congestion occurs at more distant points in the Internet, one must indirectly infer congestion. There are two recognized signals of Internet congestion: increases in latency (due to the formation of queues) and packet loss. Latency fluctuations may have many causes, complicating inference. Brandon Schlinker (Facebook / USC) investigated the inference of congestion based on passive observation of packet losses. He instrumented a sending TCP to detect lost packets, and explored congestion dynamics using a testbed that emulated a wide variety of real-life traffic conditions. They analyzed the impact of path latency, buffer size, congestion control algorithms, file size distributions, and connection reuse on congestion. In the context of Facebook traffic (typically many small transfers), they identified a simple estimator of congestion, which was the number of packets sent to first loss. Congestion comprises two stages – latency-increasing and loss-inducing. They found that an interaction between TCP and deep buffers, combined with CDN connection behavior, can yield conditions in which demand at a link can exceed the link’s capacity for an extended period of time, yet result in zero or only minimal packet loss. This research illustrates the power of passive measurement to detect and understand network conditions. However, it is probable that use of this method is restricted to the location of the sending TCP.

It is not clear that a monitor at another point along the path could make the same inference. Since the monitor observes only its own traffic, privacy concerns are side-stepped.

Narrowband IoT cellular technology. Narrowband IoT (NB-IoT) is a wide area cellular technology for connecting low power devices, with more configuration parameters and power management techniques than traditional cellular technologies. The emphasis on battery life challenges the task of performance and reliability measurements, in particular, latency, loss, and throughput can no longer capture network performance without reference to associated power consumption. Ahmed Elmokashfi (SimulaMet) gave a brief overview of this technology, presented performance results from two commercial deployments and two different NB-IoT devices, and explained how performance measurements should be taken and assessed.

5 WORKSHOP PARTICIPANTS

The main reason we continue this workshop is the enthusiastic participation it attracts from some of the brightest and most productive people in the community. We are grateful for their engagement and insights, many of which are reflected in this report.

kc claffy (UCSD), Kevin Bock and George Hughey (U. Maryland), Richard Brooks (Clemson U.), David Clark (MIT), Ann Cox (DHS S&T Cyber Security Division), Casey Deccio (BYU), Amogh Dhamdhere (Amazon Web Services) Zakir Durumeric and Liz Izhikevich (Stanford), Ahmed Elmokashfi (SIMETRIC/Simula), Avi Freedman (Kentik), Steve Huter and David Teach (U. Oregon, NSRC), Scott Jordan (UC Irvine), Ginga Kawaguti (NTT), Robert Kisteleki (RIPE NCC), Ioana Livadariu (Simula), Matthew Luckie (U. Waikato), Alexander Marder (U. Pennsylvania), Hideki Nojiri (NTT), Gregory Petropoulos (Security Scorecard), Justin Rohrer (NPS) Brandon Schlinker (Facebook, USC), Jared Smith (U. Tennessee, ORNL), Neil Spring (U. Maryland), Stephen Strowes (RIPE NCC), Renata Teixeira (INRIA, Stanford), Kevin Vermeulen (Sorbonne), Shiwei Zhang (Southern U. of S&T); and from CAIDA: kc claffy, Alberto Dainotti, Roderick Fanou, Marina Fomenkov, Alex Gamero-Garrido, Shuai Hao, Young Hyun, Alistair King, Ricky Mok, Ramakrishna Padmanabhan, and Joshua Polterock.

ACKNOWLEDGMENTS. The workshop was funded by the Department of Homeland Security (DHS) Science and Technology Directorate, Cyber Security Division (DHS S&T/CSD) contract DHS S&T contract 140D7018C0010 and the National Science Foundation OAC-1724853 and NSF CNS-1513283. The work represents the position of the authors and not necessarily that of DHS or NSF.