

On the Security of End-to-End Network Measurements

Ghassan Karame
Department of Computer Science
ETH Zurich
8092 Zurich, Switzerland
karameg@inf.ethz.ch

Srdjan Capkun
Department of Computer Science
ETH Zurich
8092 Zurich, Switzerland
capkuns@inf.ethz.ch

ABSTRACT

End-to-end network measurement tools are gaining increased importance in many Internet services. These tools were designed, however, without prior security consideration which, given the current adversarial Internet, renders their extracted network estimates questionable. In this paper, we highlight the major security vulnerabilities of existing end-to-end measurement tools and we sketch possible research directions to counter these threats. We further argue that finding “full-fledged” solutions for these challenges might require a serious re-design of the current Internet.

1. INTRODUCTION

Network measurements are becoming crucial for the operation and security of the Internet, and of several services including application-level multicast trees, content distribution and peer-to-peer (P2P) systems [1]. Numerous tools for estimating network performance have been proposed (e.g., *bandwidth measurement*: Sprobe [5], *latency*: traceroute [6], ping [8], *link quality*: mtr [7], etc.). Most of these tools push the measurement function to the end-hosts, thus conforming to the well-known end-to-end principle [2] that suggests the move of functions closer to the applications and to the end-hosts. Unfortunately, the increasing dependence of current applications and services on network measurement tools is showing the limits of foresight in the design of these tools:

- **No prior security considerations:** Current network measurement tools were developed without prior security considerations, which makes them vulnerable to *external* and *internal* attacks ranging from IP spoofing to delay and rushing attacks. Since the measurements are performed end-to-end, the end-hosts might not be able to distinguish these attacks from “authentic” measurements. These security vulnerabilities might also affect the operation of the applications that make use of these measurement tools, thus increasing the gain of the attacker.
- **Implicit trust assumptions:** Current network measurement tools were designed to work in a “friendly” environment and implicitly assume that both end-hosts are honest and behave “correctly”. However, in many situations, the interests of end-hosts might differ considerably and

thus, they might have strong incentives to cheat and increase their advantage in the network [1], [13] (e.g., free-riding). Indeed, if the endpoints misbehave and do not obey the measurement protocol, the estimated end-to-end metric will not reflect the authentic state of the network. To the best of our knowledge, the division of trust between end-hosts has not been yet analyzed in the context of network measurements.

This paper examines the security vulnerabilities of existing end-to-end network measurement tools, presents a list of design goals that *secure* measurement tools should satisfy to cope with adversarial behavior in the Internet and highlights possible research solutions in this area. Namely, we argue that finding “clean-slate” security solutions might require re-designing the architecture of the current Internet. Indirectly, our paper also motivates the need for a *secure* next-generation Internet that considers secure network measurements as first-class network citizens [21].

The remainder of the paper is organized as follows. Section 2 presents a brief outline of existing end-to-end network measurement tools. Section 3 compiles a list of security threats encountered in existing measurement tools. Section 4 suggests a list of desirable properties that secure network measurement tools should satisfy. In Section 5, we sketch possible avenues to alleviate attacks against existing tools. In Section 6, we extract relevant lessons for the future. Section 7 describes related work and we conclude in Section 8.

2. END-TO-END MEASUREMENT TOOLS

Network measurements constitute nowadays an important mean of understanding the performance and properties of network applications and systems. Network measurement tools also have the potential to solve problems in several areas such as network management, P2P architectures, etc..

2.1 Why “End-to-End” ?

Till recently, the end-to-end principle [2] has provided a justifiable rationale for moving functions closer to the end-hosts that use the function and has shaped the way the current Internet is designed. The true leverage of the end-to-end argument was implicitly a global architecture comprising a “naive” network and “smart” applications that do not require

functionality from the switching elements deployed within the network. Given this, the design of network measurements tools equally adopted the end-to-end principle, owing to the unavailability of infrastructural support for measurements and due to the absence of viable alternatives. Unfortunately, the design of current tools is revealing the limits of foresight; indeed, while the original end-to-end principle assumed a setting where both end-points are honest and behave correctly, today, users might have strong incentives to cheat [1] in order to increase their benefit in the network.

2.2 Passive Network Measurements

Passive network measurement tools rely on monitoring existing traffic between end-hosts to extract their estimates. These tools are scalable since they do not generate extra traffic for measurement purposes. Most current tools use packet capture libraries to monitor traffic and to extract various metrics of interest. Several tools for passive measurements exist, such as Nettimer [3] (for bandwidth estimation), Viznet [9] (for throughput tests), sting [10] (for latency tests), etc..

However, passive measurement tools are finding less applicability nowadays due to the fact that existing traffic is not suitable for them to produce an indicative estimate [5].

2.3 Active Network Measurements

A large number of tools *actively* generate probe traffic to extract their estimates about the network state. Most of these tools are single-ended and do not require cooperation from the measured host. Several tools for active network measurements have been proposed; CAIDA [11] hosts a comprehensive list of these tools along with their measurement functions. Examples include Pathchar [4], and Sprobe [5] measuring the *bottleneck bandwidth*, Cprobe and Spruce measuring the *available bandwidth*. Several other tools aim at measuring *the roundtrip times (RTT)*, *latency*, *delay* and *the throughput* of an end-to-end path, such as traceroute [6], mtr [7], ping [8]. A detailed description of the operation of these tools is out of the scope of this paper.

System Model: While they might be different in purpose and technique, most *active* end-to-end measurement tools share a similar system model consisting of a *verifier* and a *prover* connected by a *network* (Figure 1). The verifier wants to *measure* and *verify* the end-to-end performance of the path to the prover. The verifier actively generates probe packets destined to the prover, who appropriately echoes back its reply probe packets to the verifier. The verifier then estimates the performance of the end-to-end path to the prover by extracting and analyzing the probe packets’ arrival times depending on the measurement technique in question. For instance, Sprobe [5] makes use of TCP’s SYN and RST packets to force the remote-host to send packet-pairs replies. The verifier then estimates the bottleneck bandwidth by computing the time dispersion between the reply packet-pairs.

In this paper, we are not concerned with the detailed op-

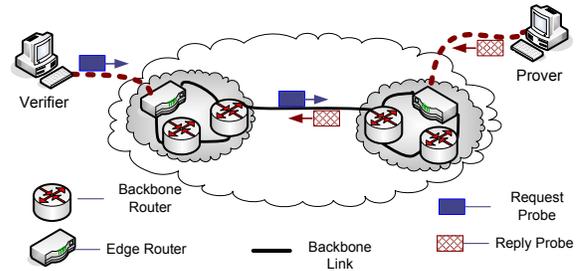


Figure 1: Active network measurements: probes are exchanged between the verifier and the prover. The former then analyzes the probes arrival/departure times.

eration of current measurement tools; instead, we abstract away their distinct features and we focus on highlighting the vulnerabilities of the general end-to-end model (illustrated in Figure 1) in adversarial settings. Given that the basic differences between these tools reside in the way they extract their estimates from the departure/arrival times of the exchanged probe packets, the model described above applies to the multitude of active end-to-end measurement tools and is not constrained by the peculiarities of a particular estimation technique. Due to lack of space, we focus on the security vulnerabilities of active measurement techniques and where appropriate we will point to the similarities and/or differences to passive measurement techniques.

3. ATTACKS AGAINST END-TO-END NETWORK MEASUREMENTS

In this section, we outline the major security threats against existing end-to-end network measurement tools.

3.1 Attacker Model

Untrusted provers constitute the core of our *internal* attacker model; untrusted provers are those hosts involved in the measurement process, but they are not trusted by the verifier to correctly execute the measurement steps. Untrusted provers can *intentionally* manipulate the sending time of their reply probes and claim a measurement value of their choice.

We also consider *external* attacks and we adopt the following attacker model: we assume that an external attacker *Eve* controls the communication channel between the verifier and the prover. Eve can also compromise routers on the path between the verifier and the prover. By compromising routers, the attacker can delay the exchanged probe packets and alter the measurements conducted by the verifier. Eve can also re-route the exchanged probes through other intermediate links to influence the conducted measurements.

3.2 Exemplary Threats

Most end-to-end measurement tools rely on simple ICMP or TCP/UDP implementations at end-hosts or at routers to exchange probes along the measured path. We identify three exemplary attacks against these tools: **delay attacks**, **rushing attacks** and **impersonation attacks**. While these attacks seem trivial and well-known, they were not addressed

in the context of end-to-end network measurements.

Delay Attacks.

Delay attacks pose a serious challenge to existing network measurements tools. An untrusted prover can intentionally *delay* its reply probes to convince the verifier of a performance value of its choice. Delay attacks can result in both inflated (higher) or deflated (lower) measurement estimates. The amount of delay Δ that needs to be introduced depends on the probe size and on the estimation techniques in use. Note that this attack is not particular to untrusted provers nor to active measurements; a *compromised router* can equally trick the verifier into accepting fake claims by introducing appropriate delays to the probes’ traversal times.

To perform delay attacks, untrusted provers can “manipulate” their networking interface and introduce appropriate delays that match their desired claims (Section 3.3). Alternatively, provers can make use of available software, such as traffic shapers (e.g., NetLimiter [15], HTB [18], etc.) to throttle their outgoing traffic according to a target rate matching their network performance claims. Delay attacks can be very hard to detect given a small unnoticeable delay. In an experiment we conducted, we were able to claim a 10 Gbps download bandwidth using the Pathchar bandwidth estimation tool [4], over a 100 Mbps physical download connection, by introducing a delay of 120 μ s.

Rushing Attacks.

An untrusted prover can also predict the *Identifier* and *Sequence Number*¹ fields in the request probe packets and send *specially crafted* reply probes ahead of time. In this way, an attacker can claim a smaller RTT which translates into a different measurement estimate.

Impersonation Attacks.

While relying on ICMP, TCP or UDP-based probe packets, current measurement tools do not provide any form of source or destination authentication. An external attacker can *spoof* the IP [14] of the prover and issue probe replies, as if it was the actual prover; the resulting metric measured by the verifier would thus be that of the attacker. This attack is further aggravated if the attacker re-routes probe packets to multiple heterogeneous workstations at its disposal to claim a performance value of her choice. Impersonation attacks can be equally mounted against passive network measurement tools; indeed, an attacker can generate traffic on behalf of a host and thus alter the network measurements.

3.3 Sample Experimental Results

To demonstrate the feasibility of delay attacks, we conducted sample experiments in which untrusted provers fake their bandwidth claims and we simulated the impact of these attacks on a typical tree-based distribution application.

¹Generally, the *Sequence Number* field in the probe requests is incremental and therefore can be easily predicted.

Since the performance of resource distribution largely depends on the organization of peers in the distribution tree (*slow* peers are the leaves of the tree and *fast* peers are located close to the multicast root), we assume that a central high-speed verifier measures the bandwidths of network hosts using the Sprobe [5] bandwidth estimation tool to optimally place them in the tree. Sprobe is a popular tool based on the packet-pair technique [12], where the verifier sends *two* back-to-back *large* TCP SYN request probe packets of equal size to the provers. The provers then issue back their TCP RST reply probe-pairs; the verifier estimates the provers’ bandwidths as follows: $B = \frac{S}{T}$, where B is the bandwidth to the prover, S is the size of the request probes and T is the time dispersion between the reply probe-pair.

Delay Attacks: Untrusted provers can claim a *higher* bandwidth by introducing a delay $\Delta = S \cdot (\frac{1}{B_{claimed}} - \frac{1}{B_{auth}})$ before responding to the first request packet. Here, $B_{claimed}$ denotes the fake claimed bandwidth of the prover and B_{auth} is the genuine bandwidth of the prover. By claiming higher bandwidths, provers are likely to be placed close to the content multicast root, and are therefore likely to acquire the resource faster. Alternatively, provers might claim *lower* bandwidths, to reduce their upload contribution in the network. Provers can claim lower bandwidths by delaying their reply to the second request probe issued by the verifier. These attacks can be equally performed by a *rogue* router located on the path between the end-hosts.

To introduce the aforementioned delays, we devised an application that replaces the kernel’s TCP/IP stack by a raw socket and uses an iptable rule to drop all replies issued by the kernel; the prover then sends back the TCP SYN reply probes with the desired delay. We conducted our measurements on 100 Mbps symmetric physical connections deployed on two different paths: **Path1** where both the verifier and the prover hosts are located in the same state, and **Path2** where the verifier is located in Europe and the prover is located in the US. Our findings are depicted in Figure 2. In the figure, *target bandwidth* refers to the bandwidth that an untrusted prover claims and *measured bandwidth* denotes the bottleneck bandwidth estimate extracted by the verifier. Indeed, as shown in Figure 2, the prover can, by appropriately delaying its reply probes, claim a bandwidth of its choice irrespective of its physical bandwidth capability.

Implications of Attacks: To study the impact of these attacks, we implemented a C-based simulator that simulates content distribution to a large number of nodes organized in a binary tree topology. In our simulations, the nodes’ bandwidth were chosen based on the bandwidth distribution in current P2P networks as reported in [1]. Each data point is averaged over 1000 runs. Figure 3 shows the detrimental impact of fake bandwidth claims on resource distribution; the average download times over *all* peers in the network *doubles* when 40 % of the peers claim a lower bandwidth than

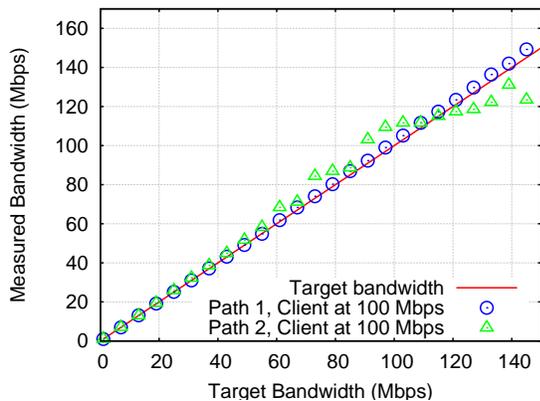


Figure 2: Delay Attacks on Sprobe [5].

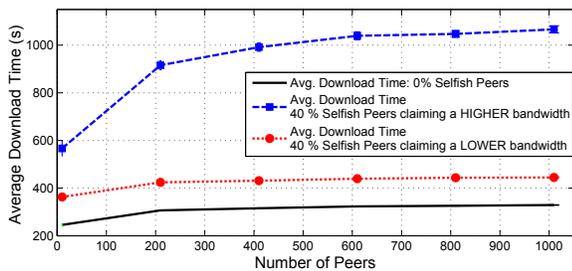


Figure 3: Impact of attacks on the peers’ average download times in multicast binary trees.

their own. This effect is more detrimental when those peers claim a higher bandwidth than they actually have; the average download time over all peers almost *quadruples*.

4. DESIGN GOALS FOR SECURE NETWORK MEASUREMENT TOOLS

In this section, we outline desirable properties that a *secure* end-to-end network measurement tool should satisfy to be widely deployed in the current Internet.

Active and Cooperative Measurements.

Some previous work [5] argues that network measurement tools should be designed to work in uncooperative environments to scale to a large number of Internet hosts without the need for dedicated software on the prover’s side. Although this is indeed a desirable property, we find support for uncooperative environments rather unrealistic. In fact, with the increasing proliferation of “de-facto” security applications, such as home firewalls, uncooperative probing techniques find less applicability in the near future as they are likely to be considered hostile by end-hosts. Some routers already *filter* ICMP packets due to their potential malevolent use [12]. We believe that support for cooperative measurement environments is inevitable in the near-future. Furthermore, end-to-end security would impair the use of measurement tools in passive and uncooperative environments as it requires end-host cooperation in achieving essential security primitives (e.g., source authentication). Unless functionality is pushed to the Internet backbone (Section 6), cooperative environ-

ments present themselves as vital “playgrounds” for secure end-to-end network measurements in the current Internet.

Resilience to External Attackers.

Secure end-to-end measurement tools should be resilient to *external attacks* on the measurement process. Since measurements are performed end-to-end, the impact of an external entity manipulating the propagation times of the probe packets can be detrimental.

Resilience to Internal Attackers.

Current end-to-end network measurement techniques push trust towards end-hosts; an untrusted prover can abuse this trust in numerous ways and claim a performance capability that increases its advantage. Secure end-to-end measurement tools should *detect* those internal attackers that are not correctly executing the measurement phases.

Efficiency and Accuracy.

Secure end-to-end measurement tools should preserve the accuracy of the measurements and cope with inherent network properties such as cross-traffic and traffic dynamism. We further argue that the whole process should be as efficient as possible. This includes the overhead in authenticating the exchanged probes, and detecting possible host misbehavior.

5. POSSIBLE COUNTERMEASURES

In what follows, we sketch potential solutions to mitigate attacks against end-to-end network measurement tools.

Countering Impersonation Attacks.

Network measurement tools can use standard authentication protocols (e.g., lightweight symmetric cryptography) to counter impersonation attacks. By authenticating the exchanged probes (using pre-shared keys for example), an external attacker will not be able to impersonate the prover by issuing probe packets on its behalf since it cannot forge the prover’s authentication on the verifier’s request probes.

Countering Rushing Attacks.

To counter rushing attacks, the verifier can use pseudo-random functions to generate its request probes such that they cannot be predicted by the provers and require that the reply probes issued by the provers are correlated in content to its request probes. For instance, the verifier can make use of distance bounding protocols [16] to ensure that the prover receives the request probes before issuing its reply. Alternatively, it can require that the prover authenticates the received pseudo-random probes using a shared key. Thus, the probability that an untrusted prover sends a correct reply probe ahead of time can be made satisfactorily negligible.

Countering Delay Attacks.

As far as we are aware, there are no means for the verifier to remotely verify that the provers did not intentionally

introduce delays before issuing their reply probes. In what follows, we highlight “best-effort” avenues to address the challenges posed by delay attacks.

Analyzing “Reference” Data: In theory, delay attacks can be alleviated if the verifier knows an estimate of the RTT to the prover. The verifier can acquire RTT references through offline measurements or from online servers that perform RTT measurements around the globe [17]. However, given the variability of RTTs in the Internet and the impact of small delays on the measurement process, this technique is unlikely to prevent untrusted provers from claiming a fake measurement value. Nevertheless, we argue that it *limits* the range of fake claims that the provers are able to make.

Alternatively, the verifier could take advantage of the reference data that it possesses to construct an elaborate opinion about the network performance to the prover, or to correlate different metrics (e.g., RTT and bandwidth). For example, if the verifier knows that the download bottleneck bandwidth to the prover is in the order of few hundreds of Kbps, it is highly unlikely that the path to the prover will have a 100 Mbps upload bottleneck bandwidth.

“Host Fingerprinting”: One viable method to filter in-authentic measurements is to acquire relevant information about the prover’s host properties, such as its operating system, type of networking interface, location, etc. The literature includes several proposals for remote operating system fingerprinting [22] and remote device fingerprinting [23]. Using such information, the verifier might detect fake claims by untrusted provers. For example, if the verifier knows that the prover is equipped with an 11 Mbps wireless card, then it is unlikely that it has a 50 Mbps download bandwidth.

Similarly, the verifier can deduce the prover’s location and its type of Internet connection from its IP address; if the IP of the prover corresponds to an IP in the same state where the verifier is located, then it is improbable that its RTT to the verifier is 300 ms. In addition, by resolving a prover’s IP address into its Domain Name Server (DNS), the verifier can detect discrepancies in the measurements; for instance, if the verifier measures a 5 Mbps download bandwidth while the prover’s DNS name is “user.dialup.com”, then it is highly probable that there was an attack on the measurements.

Detecting Traffic and Bandwidth Shapers: Traffic shapers (e.g., NetLimiter [15]) present themselves as effortless and powerful routines to conduct delay attacks on end-to-end network measurement tools. Traffic shapers limit the amount of data a host transmits and/or accepts by delaying incoming and/or outgoing packets to match a specified rate input by users. Our findings show that the rate at which a small burst of packets can be shaped is dependent on the kernel’s system timer frequency. This might, therefore, give a considerable advantage for the verifier to detect their deployment at the prover’s side. We are currently investigating the feasibility of remote detection mechanisms of traffic shapers.

Outlier Detection and Reputation-Based Schemes: Reputation-based approaches [19] and statistical outlier detection [13] can also be used to prevent untrusted hosts from tampering with the measurement process. For instance, a dedicated server can keep *history* of the recorded measurements. Using statistical outlier detection methods, fake claims can be detected and malicious hosts can be identified. Similarly, each host can be associated with a *reputation value* that indicates how trustworthy it is. Hosts that vary their claims will be associated with low reputation values and, therefore, will not be chosen in subsequent interactions.

Tamper-Resistant Network Interface: As far as we aware, it is very hard, if not impossible, for the verifier to remotely detect tampering with the networking interface of the provers. One solution that addresses this problem is to embed networking tamper-resistant software or hardware on the provers’ side to prevent them from tampering with their network interface. Unfortunately, tamper-proof software/hardware come at high implementation costs nowadays. Given that they will be used only during the network measurement process, we do acknowledge, therefore, that relying on tamper-resistant software/hardware might constitute a rather “bulky” proposal.

6. OUTLOOK

In this paper, we highlighted exemplary attacks on network measurement tools and we suggested “best-effort” techniques to mitigate these attacks. However, these techniques fail to fully eliminate the multitude of potential threats against existing tools. Namely, mitigating delay attacks on measurements remains, to a large extent, an open research problem. We believe that the easy realization of these attacks renders the estimates of current tools highly questionable. Surprisingly, while the evolution of the Internet stimulated a serious re-interpretation of the end-to-end principle, the design of end-to-end measurement tools remains unaddressed and lacks basic security primitives, such as *trust-to-trust* factors.

Given the current trends in designing a “clean-slate” future Internet, our paper also motivates the need for a *secure* next-generation Internet. We argue that any next-generation Internet design must provide robust functionality to support secure network measurements. This means that network measurements should be considered in all aspects of design, deployment and operation. Since network measurements are gaining paramount importance in monitoring the performance of the Internet, *secure* infrastructural support for network measurements becomes rather a necessity. We can identify a number of reasons why it might be beneficial to embed parts of the measurement functions “inside the network”: 1) by performing measurements at an intermediate point, end-users can avoid the cost and overhead of generating unwelcome traffic across the network, 2) by pushing functionality from end-hosts to *dedicated* and *trusted* network components, several security threats can be eliminated. For instance, edge routers could securely timestamp incoming and

outgoing probes and identify whether queuing has occurred. This would facilitate the detection of delay attacks. Performance “awareness” is another desirable design property for next-generation Internet. Dedicated network components could in the future construct and store bandwidth and latency maps of Internet hosts by monitoring incoming and outgoing traffic. This would indeed eliminate the need for *insecure* and “bulky” end-to-end measurement tools.

However, even if these challenges are overcome through an infrastructure that supports network measurements, we still expect other hazards to arise with respect to the design of future measurement components. Indeed, since the Internet traffic pattern is becoming increasingly dynamic, the underlying measurement infrastructure should be flexible to accommodate for such dynamism without the need of frequent re-design. Furthermore, the use of dedicated measurement components requires the presence of trusted authorities that control and maintain these components. Here, delicate access control to the infrastructure is a crucial design property in order to address privacy and integrity concerns; any compromise of the measurement components might result in severe performance deterioration of the entire network. Overcoming these limitations requires a very careful design of the measurement functions in the future Internet. We therefore hope that our paper motivates further research in this area.

7. RELATED WORK

Several tools for network measurements have been proposed and evaluated empirically over a number of Internet paths. CAIDA [11] hosts a comprehensive list of these tools along with their measurement functions. Examples include: Nettimer [3], Sprobe [5], traceroute [6], mtr [7], ping [8], etc.. However, none of these tools consider security issues in end-to-end network measurements. For instance, Sariou *et al.* suggest in [5] ideal design goals for bandwidth estimation tools, but implicitly assume a setting where end-hosts are honest and abide by the measurement protocol.

Although few schemes were proposed in the scope of securing RTT measurements [20], they assume a context where the prover attests that it is within the proximity of the verifier; the prover, therefore, does not have incentives to mount delay-based attacks in this setting. This is not the case in current Internet applications; an untrusted prover might have considerable incentives to fake its performance metrics to the verifier [1]. In [13], Walters *et al.* propose to mitigate attacks against measurements in overlay networks by combining anomaly detection and reputation-based techniques.

8. SUMMARY

In this paper, we highlighted the major security vulnerabilities in existing measurement tools and we sketched possible avenues to counter these threats. Another aim of this paper is to increase the awareness in the community of the design issues for these tools. As far as we are aware, no previous contribution has analyzed the security implications of

the design of network measurement tools. Due to the lack of foresight in their design, it is very hard, if not impossible, to fully counter all security challenges against existing tools. More specifically, delay attacks pose serious challenges to the consistency of network measurements as they constitute simple, yet powerful, means for attackers to abuse the operation of existing tools and maximize their profit in the network. Full-fledged countermeasures might require the design of a “clean-slate” future Internet that provides trusted infrastructure for secure network measurements.

9. REFERENCES

- [1] S. Sariou, P. Gummadi, S. Gribble, “A Measurement Study of Peer-to-Peer File Sharing Systems”, In *MMCN* 2002.
- [2] J. H. Saltzer, D. P. Reed and D. D. Clark, “End-to-End Arguments in System Design”, In *ACM Transactions on Computer Systems*, 1984.
- [3] K. Lai and M. Baker, “Nettimer: A Tool for Measuring Bottleneck Link Bandwidth”, In *USITS*, 2001.
- [4] V. Jacobson, “Pathchar”, <http://www.caida.org/tools/utilities/others/pathchar>, 1997.
- [5] S. Sariou, P. Gummadi and S. Gribble, “SProbe: A Fast Technique for Measuring Bottleneck Bandwidth in Uncooperative Environments”, In *INFOCOM*, 2002.
- [6] “Traceroute”, <ftp://ftp.ee.lbl.gov/traceroute.tar.gz>.
- [7] “mtr”, <http://www.bitwizard.nl/mtr/>.
- [8] M. Muuss, “ping”, <ftp://ftp.arl.mil/pub/ping.shar>.
- [9] “Viznet”, <http://dast.nlanr.net/Projects/Viznet/>
- [10] S. Savage, “Sting”, <http://www.cs.ucsd.edu/~savage/sting/>
- [11] CAIDA, tools: taxonomies, <http://www.caida.org/tools/taxonomy/performance.xml>.
- [12] K. Lai, M. Baker, “Measuring link bandwidths using a deterministic model of packet delays”, In *SIGCOMM* 2000.
- [13] A. Walters, D. Zage and C. Nita-Rotaru, “A Framework for Mitigating Attacks Against Measurement-Based Adaptation Mechanisms in Unstructured Multicast Overlay Networks”, In *ACM/IEEE Trans. on Networking*, 2007.
- [14] B. Harris and R. Hunt, “TCP/IP security threats and attack methods”, In *Computer Communications*, 1999.
- [15] NetLimiter, <http://www.netlimiter.com/>.
- [16] S. Brands and D. Chaum, “Distance-bounding protocols”, In *EUROCRYPT*, 1994.
- [17] The CAIDA DNS root/gTLD RTT Dataset, <https://data.caida.org/datasets/dns/root-gtld-rtt/>.
- [18] HTB Traffic Shaper, <http://luxik.cdi.cz/~devik/qos/htb/>.
- [19] S. D. Kamvar, M. T. Schlosser, H. Garcia-Molina, “The EigenTrust Algorithm for Reputation Management in P2P Networks”, In *WWW* 2003.
- [20] O. Courtaf, M. Karroum and A. Duran, “Method and Devices for Secure Measurements of Time-Based Distance Between Two Devices”, *Patent no. WO/2006/136278*, 2006.
- [21] P. Barford, “Measurement as a First Class Network Citizen”, <http://pages.cs.wisc.edu/~pbarford/sngi-whitepaper.pdf>.
- [22] Nmap free security scanner, 2004, <http://www.insecure.org/nmap/>.
- [23] T. Kohno, A. Broido, and K.C. Claffy, “Remote physical device fingerprinting”, In *IEEE TDSC*, 2005.