

On the Performance of Round Trip Time Network Tomography

Yolanda Tsang, Mehmet Yildiz, Paul Barford, Robert Nowak.

Abstract—Network tomography is an appealing method for active measurement of link level characteristics such as delay and loss on end-to-end paths. Most network tomography techniques developed to date are based on *one-way* measurements requiring collaboration from both sending and receiving hosts which severely limits the scope of the paths over which these techniques can be used. We extend our previous work on *Network Radar*, a new tomographic inference method based on round trip time (RTT) measurements from TCP SYN/SYN-ACK packets. In this paper, our contributions are three-folded. (1) We extend our analytic framework for estimating delay variance on the shared network segment using Network Radar to include confidence estimates which enable measurement accuracy to be assessed - an important consideration for practical deployment. (2) We evaluate Network Radar in a series of experiments conducted in a controlled laboratory environment. These tests explore the boundaries of effectiveness of our RTT-based method, and show that it works well over a wide range of traffic conditions. (3) We evaluate Network Radar in a series of tests conducted in the wide area Internet. These tests show that RTT-based delay variance estimates can be used effectively to identify most likely network topology - a natural and verifiable application for RTT tomography. The performance results in this paper demonstrate that Network Radar can now be used for both research and operational purposes.

I. INTRODUCTION

Researchers interested in measuring the behavioral and structural characteristics of the Internet face many challenges, not the least of which is lack of open instrumentation in the infrastructure. Similarly, network operators charged with the responsibility of identifying and diagnosing problems in their networks have an on-going need for measurement tools that make this process faster and more efficient. These issues have led to the development of novel methods and systems for measuring network characteristics beyond the confines of a single network. These tools report network characteristics based on either passive measurements of traffic at a given vantage point, or by measuring the response to probe packets emitted by the tool itself.

One class of probe-based measurement tools developed over the past several years are those that use *tomographic inference* techniques to deduce link-specific information such as packet loss rates, packet delay characteristics and network topology. Prior work has established the basic mechanisms for the use of tomographic inference in the networking context [1], [2], [3], [4]. Despite these interesting results, validation and wide-spread deployment of tools based on these techniques is problematic. Most of these network tomography tools are based on the use of *one way* probe traffic measurements, and as such, they require *cooperation between sending and receiving hosts* (the only exceptions to this requirement that we are aware of are methods based on passive monitoring loss [5], [6], [3]). This requirement places a significant limit on the scope of the paths over which measurements can be made, and thereby limits wide-spread use

of these tools for research and/or network management activities. An additional practical limitation of one-way tomographic tools that measure delay statistics is that they require clock synchronization between end hosts. While new synchronization methods such as those proposed in [7], [8] offer promise, synchronization of large numbers of end host clocks may not happen for some time.

Recently, a new network tomographic technique based on round trip time (RTT) measurements was proposed which eliminates the need for special-purpose cooperation from receivers [9]. This technique, known as *Network Radar*, uses RTT measurements from TCP SYN and SYN-ACK segments to estimate the delay variance of the shared network segment in a standard one sender–two receivers configuration. In this work, the authors evaluated Network Radar in a series of preliminary tests conducted in an emulation environment. Although promising, it is difficult to access the effectiveness of Network Radar in the Internet.

In this paper, we investigated the conditions under which Network Radar is effective by careful designing a series of laboratory-based and Internet-based experiments. The emulation experiments were conducted on a network of Cisco routers and PC hosts configured in a simple topology. The experiments revealed that Network Radar is effective over a wide range of traffic conditions. Most importantly, using our proposed confidence estimator, we were able to establish a correlation between the delay estimator and the number of measurements.

Our experiments in the Internet examine the application of RTT tomography to the problem of logical topology discovery. In this case, we use a clustering algorithm [1] based on measured shared segment delays to establish logical connectivity between nodes, but unlike [1], our RTT measurements are more prone to the variability of delays. We choose this application since it gives us an opportunity to validate our results using other tools whereas delay or loss tomography would have required instrumentation that was not readily available in the wide area. Our experiments consist of randomly selecting a set of destinations and then using RTT tomography to discern topological connectivity between a source and those destinations. We validated the results using *traceroute*. The results show that RTT tomography is very effective at establishing logical topology.

In summary, this paper makes the following contributions. First, we extend the analytical framework based on RTT measurements to include an automatic confidence estimator that enables measurements that would otherwise skew results to be discarded. This is an important capability for practical use of Net-

Y.Tsang and R. Nowak are with the ECE Dept. of University of Wisconsin – Madison. E-mail: ytsang@ieee.org and nowak@engr.wisc.edu. M.Yildiz is with the ECE Dept. of Cornell University. E-mail: mey7@cornell.edu. P.Barford is with the CS Dept. of University of Wisconsin – Madison. E-mail: pb@cs.wisc.edu.

work Radar. Second, we investigate the robustness of RTT tomography in a series of controlled laboratory experiments and show that it is effective over a wide range of operating conditions. Third, we investigate the capability of Network Radar in Internet-based tests that we can validate with supporting measurements and demonstrate that our tool is effective in the wide area.

II. RELATED WORK

Network tomography has been an emerging area of study in the past few years. A variety of inference techniques in estimating loss rates, delay statistics and topology identification have been treated in [1], [2], [10], [11], [12]. Most of these techniques are promising but are not widely applicable because of the requirement to have cooperation from receivers to collect one-way measurement. Most of these techniques also require clock synchronization between end hosts or they depend on special cooperation from the internal routers. These practical limitations have largely prevented the tomographic techniques from being used widely.

The idea of RTT-based network tomography is introduced in [9]. In that paper, we develop an analytic foundation for this technique and discuss preliminary considerations for widespread use which we treat in this paper. There are other RTT-based measurement studies in inferring path characteristics. Unlike ours, they are based on either the time-to-live (TTL) [12] or the timestamp [10] options in the Internet Control Message Protocol (ICMP). Security and privacy concerns related to the use of ICMP options have led to both rate limiting and outright blocking of ICMP packets, thus reducing the effectiveness of ICMP-based tomography. Our TCP-based methodology is easy to deploy for asymmetric path and is widely available.

III. METHODOLOGY AND MEASUREMENT FRAMEWORK

In this section, we briefly introduce the basic concept of network tomography. Traditional network tomography assumes a single source transmitting probe packets to pairs of receivers. The topology is assumed to be fixed throughout the measurement period (*i.e.*, the routing table does not change and no load balancing¹ is employed) forming a tree with the *source* at the root and the *receivers* at the leaves in Figure 2. The branching node between the source and receivers represents an internal *router*. Connections between the source, router, and receivers are called *segments* or logical links. Each segment between may be a direct connection, or there may be “hidden” routers or switches (where no branching occurs) along the path that are not explicitly shown. We focus specifically on *delay variance estimation* as in [9] on shared segment of the path. This focus is by convention only and there are nothing inherent in our descriptions that prevent the tomography methods from loss rate measurements.

Basic tomographic measurement and inference ideas are straightforward. Assume that the individual link delays along

¹Load balancing is typically prefix based in the current Internet, as contrast to load based, to avoid packet reordering. Thus the packets follow the same path between a source-destination pair.




Fig. 1. Round trip time tomographic delay variance estimation in a standard one sender (0) two-receiver (1, 2) network. Variances on shared (σ_s^2) and unshared segments (σ_1^2, σ_2^2) are noted.

the end-to-end path are independent and stationary. A sender transmits two closely time-spaced (back-to-back) probe packets to different receivers. The paths to these receivers traverse a common set of links, but at some point the two paths diverge (as the tree branches). The two packets are expected to experience approximately the same delay on the shared segment in their path. The round trip delay consists of transmission, propagation, processing and queueing delays. The delay variances are mainly caused by the queueing delay, where the rest of the delays can be modeled as a nearly constant quantities. Because the delays on the shared and unshared links are assumed to be independent, a straight-forward calculation shows that the delay variances on each path equals to the sum of delay variances on each link. The above observation has been used in topology identification. This also allow us to evaluate and to demonstrate the performance of the Network Radar, which will be detailed in Section IV-B.

Network Radar measures round trip times by sending TCP SYN packets to the HTTP service (port 80) on a target end host. Any remote host running this service will respond with a SYN-ACK packet. Round trip time measurements can then be made at the sender using simple time differencing between transmission of the SYN and receipt of the SYN-ACK. In cases where a port does not have any associated HTTP server, most hosts will send back a RST when receiving a SYN.

IV. PROPOSED ESTIMATOR AND PERFORMANCE ANALYSIS

In this section we present an unbiased estimator for the shared link variance, σ_s^2 , and derive confidence intervals for the estimator. The confidence intervals allow us to automatically detect and reject cases in which the estimator is unreliable. The standard result of an unbiased estimator for the variance can be found in most statistics textbook. However, in this section, we derive the explicit formula for the unbiased estimator given the conditions and assumptions in our context.

A. Confidence Estimator

To begin our analysis, we formally demonstrate that, under the stated assumptions, the covariance of the RTTs is equal to the delay variance on the shared link.

Proposition 1: Denote the N RTT packet pair measurements by $\mathbf{y} \equiv \{y_1(k), y_2(k)\}_{k=1}^N$, where $y_1(k)$ and $y_2(k)$ are the k th RTT measurements to/from receiver 1 and 2 respectively. The delay on the shared path is denoted by $d_s(k)$, and the de-

lays on the unshared paths are $d_1(k)$ and $d_2(k)$, where k denotes k th RTT pair. Assume that the packets in the k -th pair experience an identical delay, $d_s(k)$, on the shared portion of their paths, and that $d_s(k)$ is statistically independent of the remainder of the RTT delays, $d_1(k) = y_1(k) - d_s(k)$ and $d_2(k) = y_2(k) - d_s(k)$. Furthermore, assume that the delays in different measurements (different k) are statistically independent. Then $\text{cov}(y_1, y_2) = \sigma_s^2$. Moreover, $\widehat{\sigma}_s^2$ is an unbiased estimator of σ_s^2 .

$$\widehat{\sigma}_s^2 \equiv \frac{1}{N-1} \sum_{k=1}^N (y_1(k) - \bar{y}_1)(y_2(k) - \bar{y}_2) \quad (1)$$

where \bar{y}_i is the sample mean of $\{y_i(k)\}_{k=1}^N$ for $i = 1, 2$.

The assumptions of Proposition 1 are approximately met in actual practice. The first assumption is that the two back-to-back packets experience the same delay on the shared path. This is reasonable due to the back-to-back nature of the probes, and we experimentally verify this assumption later in the paper. The second assumption regarding the statistical independence of the delays is reasonable because the cross-traffic tends to be independent on the shared and unshared paths, and by sufficiently spacing the probes in time, the delays in different measurements are fairly independent. Both these assumptions have been verified by experimental work as well. Also, the assumption of independence could be weakened to only assume that the delays are uncorrelated.

Next we investigate the reliability of the estimator $\widehat{\sigma}_s^2$ via confidence intervals. Specifically, we first determine the theoretical variance of the estimator, and then propose an unbiased estimator for this variance. The square-root of the variance equals the standard deviation δ of $\widehat{\sigma}_s^2$, which provides confidence intervals of the form $\widehat{\sigma}_s^2 \pm \alpha \delta$, where $\alpha > 0$ is the confidence level (e.g., $\alpha = 3$ produces a confidence interval of three standard deviations).

In practice, we do not have the theoretical value of the standard deviation δ , but we can estimate it from the data as shown in the next proposition.

Proposition 2: Define

$$\widehat{\delta}^2 \equiv \frac{1}{(N-3)(N-2)} \sum_{k=1}^N (y_1(k) - \bar{y}_1)^2 (y_2(k) - \bar{y}_2)^2 - \frac{N+1}{N(N-3)} (\widehat{\sigma}_s^2)^2 - \frac{N-1}{N(N-2)(N-3)} \widehat{\sigma}_1^2 \widehat{\sigma}_2^2 \quad (2)$$

where $\widehat{\sigma}_i^2$ is the empirical variance of the RTTs to destination i , for $i = 1, 2$. Then under the assumptions of Proposition 1, $E[\widehat{\delta}^2] = \delta^2$; i.e., $\widehat{\delta}^2$ is an unbiased estimate of the theoretical value δ^2 .

Finally, we apply the confidence intervals to automatically detect and reject cases in which the estimator $\widehat{\sigma}_s^2$ is unreliable. From the propositions, and the measured data, we can compute $\widehat{\sigma}_s^2$ and a confidence interval of the $\widehat{\sigma}_s^2 \pm \alpha \widehat{\delta}$, where $\widehat{\delta} \equiv (\widehat{\delta}^2)^{1/2}$. If $\widehat{\delta} \gg \widehat{\sigma}_s^2$, then the confidence interval is quite large relative

to the center point. In such cases, we can be reasonably conclude that the estimator is unreliable. The unreliability could be due to excessive “noise” on the unshared portions (as mentioned above) or large deviations from the assumptions of our the theory. Thus, for a given confidence level α , we say that the estimator is α -unreliable if

$$\frac{\widehat{\sigma}_s^2}{\widehat{\delta}} < \alpha.$$

In practice, we recommend a level of $\alpha \geq 3$ where α is also known as the Signal to Noise Ratio (SNR), the ratio between the estimator and its confidence.

B. Topology Identification

As mentioned earlier, the topology is tree-structured. We can thus apply Network Radar to pairs of receivers to localize the delay variance of each individual (logical) link. This feature also allows us to use topology identification as a *verifiable* method for demonstrating the performance of the RTT tomographic tool. Indeed, one of the significant challenges of network tomography studies in general is *in situ* evaluation since access to network internal instrumentation is so rarely available. In the context of topology discovery, end-to-end topology measurement tools such as traceroute or Scriptroute [13] which can be used to validate the logical topology structures generated using the aforementioned RTT tomography and clustering methods described next.




Fig. 2. An example of topology identification illustration with 3 receiver nodes.

An assumption in standard network tomography as well as our framework is that the delays on the shared and unshared links are independent. Therefore, the end-to-end delay variance will be the sum of the link delay variances along the path. One can also use this information to identify the underlying network topology configuration. This is done by interpreting estimates of delay variation as a measure of the physical extent of the shared path. Let p_{ij} denote the shared path between the source and a pair of receivers i and j . Let i, j and k be three receivers and suppose that p_{ij} is longer (from a physical distance perspective) than p_{ik} (clearly one of the shared paths has to be included in the other). The delay variance associated with p_{ij} (denoted $\sigma_{p_{ij}}^2$) is going to be larger than the delay variance associated with p_{ik} . This partial ordering of the delay variance of the shared paths can be used to identify the underlying topology. For example, refer to Figure 2, $\sigma_{p_{1,2}}^2$ will be greater than $\sigma_{p_{i,2}}^2$ for any other receiver i , revealing that receivers 1 and 2 are siblings in the logical tree. This property can be exploited to devise a simple and effective bottom-up merging algorithms that identifies the full, logical topology from the pairwise delay variances. This problem is indeed one of hierarchical clustering.

In general we do not have access to the exact delay variances, but only noisy measurements of those. In particular, the variability of the measurements for each pair of receivers may be significant. In [1] a hierarchical clustering technique that is suitable for the topology identification problem is proposed. This technique is based on a likelihood model capable of accounting for the different statistical properties of the measurements for each pair of receivers. Two clustering techniques building on the likelihood model are proposed in [1]: (i) a deterministic, greedy approach, with very low computation complexity; (ii) a Markov Chain Monte Carlo method that attempts to find a globally optimal (in a likelihood sense) topology, at the expense of computation power. In the Section V-B, we apply these techniques to infer the topology solely from RTT covariance measurements.

V. EXPERIMENTAL EVALUATION

We evaluate the capabilities and robustness of RTT tomography in a series of lab-based and Internet-based experiments. The lab-based experiments enable control and instrumentation of all aspects of the test environment. This allows us to verify the capabilities and limitations of the tool in a realistic but limited infrastructure. We also perform a series of Internet experiments to assess the tool's performance in the wide area network. Its capability in identifying the logical network topology provides an indirect but verifiable (via *traceroute*) means of assessing the tool's performance. A series of more extensive tests are also available in [14].




Fig. 3. The laboratory network configuration includes 5 routers and 9 PCs. The sending host is 0 and the receivers are 1 and 2 (logical topology in gray). The boxes xT denote cross-traffic generators and the balls R denote routers. $S1$ and $S2$ denote measurement systems used to validate the performance of the RTT based tool and DAG denotes the DAG measurement system placement.

A. Emulation Experiments

Our lab-based experimental environment carried out in the Wisconsin Advanced Internet Laboratory [15] includes 5 Cisco commercial routers (7200/7500/12000 series) and 9 PCs running Redhat Linux. The bandwidth on all connections is $1Gb/s$. The setup is illustrated in Fig. 3. Boxes 0, 1 and 2 denote the nodes of interests as in Fig. 1. Box 0 refers to the sender and box 1 and 2 are the receivers. Variable background (non-probe) traffic in this environment is generated using *Harpoon* [16], a flow level traffic generator that runs on the systems denoted by xT . Propagation delays on individual links are emulated using a simple configuration of the Click modular router [17]. During each experiment, background traffic loads are generated based on input distributions derived from NetFlow logs captured at the border router of University of Wisconsin. Emulated propa-

gation delays on each link are fixed and remain constant.

Each measurement period consist of 1000 packet pairs sent from node 0 (the sender) to receiver nodes 1 and 2. The send rate is fixed at a rate of $10probes/sec$ (100ms intervals). Practically speaking, this is a low probe rate which should not cause concern for medium to large web sites in the Internet. At the end of each measurement period, we collect packet traces from *tcpdump* which is running on the sender (node 0) and at two monitoring devices ($S1$ and $S2$) along the path to the receivers. The monitors, which of course are not possible outside of the lab, allow us to verify the performance of our tool by providing ground truth measurements of packet delays. The first monitor, $S1$, records the back-to-back packet spacing entering the branching router. The second monitor, $S2$, records outgoing packets from the branching router 2 with extra cross traffic and it provides us the “true” delay variance on the shared link of the path. We synchronize the clocks on the monitoring hosts via network time protocol. The clocks are disciplined from a local stratum 1 time source giving us synchronization on the order of single milliseconds. This is important for verification of one-way delays in our experiments.

A.1 Results

Figure 4 depicts the accuracy of our tool, Network Radar, by comparing the estimates to the “true” value of the shared path delay standard deviation. The “true” value (σ_s) for the one way delay on the shared path is the measured time difference of TCP SYN packets at the sender and at the second monitor $S2$ in Figure 3. The estimates are computed directly from y_1 and y_2 , the end-to-end measurements to receiver. A moderate level of background traffic was used during this experiment and no packet loss was observed. The one-way delay between node 0 and node 1 was fixed at 0.6ms while the node 0 to node 2 delay was fixed at 6ms. Ideally, the estimates should be identical to the “true” value and fall onto the 45° line. We hypothesize that the discrepancy might arise if the timestamping mechanism is inaccurate or unreliable, or if the back-to-back assumptions described in the previous section are violated. Nonetheless, the estimates are close to the real value.




Fig. 4. Plot of standard deviation from direct measurement (horizontal axis) vs. estimated delay standard deviation using RTT-tomography tool (vertical axis).

A.2 Operating Conditions Study

One of the primary objectives in our lab-based experiments is to understand the operating conditions under which RTT tomography is practical. To investigate potential sources of er-

rors that would make RTT tomography ineffective, we examine the effects of; (1) the reliability of the timestamping mechanism (comparing `tcpdump` and others), (2) varying levels of background traffic, (3) varying spacing between packet pairs as they are emitted by the sender (back-to-backness), (4) load on receiving end-hosts which can cause variability in generation of response packets (SYN-ACK generation delay), and (5) number of measurements as compared to the resulting estimation error. Except the case where we investigate the effects of the background traffic, we fix the amount of background traffic load.




Fig. 5. Verification of timestamping mechanism using `tcpdump` utility and hardware-base reference systems using DAG cards.

Timestamping mechanism We evaluate our timestamping mechanism with the Endace DAG cards [18], the current state of art in hardware timestamping. It provides timestamping with high accuracy, with no packet loss and no extra delay. Most network interface cards suffer from corrupted, delayed or lost packets when the traffic volume is high. We use systems with DAG cards to assess how `tcpdump` timestamping variability at the sender PC effects RTT measurements. During these experiments we capture the packet departure and arrival times using both `tcpdump` and systems with DAG cards. By comparing the results in Fig. 5, we conclude that the variability in RTT measurements caused by `tcpdump` does not have a significant impact. For the rest of the paper, we collect the measurements with `tcpdump` timestamping mechanism.




Fig. 6. Significant of background traffic in the estimates. High SNR symbolizes high accuracy in the estimates. The signal to noise ratio (SNR) increases when the background traffic (correlation) on the shared segment increases. Unshared segments had moderate load in these experiments.

Background Traffic With respect to background traffic, the capabilities of RTT tomography are proportionally dependent on the load on the shared segment and inversely proportional to the load on the unshared segments. By fixing the traffic on

HTTP Connections	Average Delay	Delay Variance
1	$7.23\mu s$	$1.8 \times 10^{-12} s^2$
201	$28\mu s$	$3.4 \times 10^{-11} s^2$

TABLE I

AVERAGE SYN-ACK GENERATION DELAY AND ITS VARIANCE WITH VARYING END HOST LOAD

the unshared segment (noise) and varying the the traffic on the shared segment (signal), it is shown in Figure 6 that the signal quality improves. Note that the accuracy of our estimator not only depends on the variance on the shared segment, but also the variances on the unshared segments. As these variabilities increase, our confidence weakens, as shown in Equation 2.




Fig. 7. Plot of directly measured delay standard deviation (vertical axis) and the estimated one (horizontal axis) on the shared link when the spacing between the packets are $3\mu s$ (circle) and $100ms$ (triangle) apart at the sender.

Back-to-back assumption Figure 7 depicts the accuracy of the results by varying the packet spacing within the probe pair while the spacing across pairs remains unchanged. The results agree with our expectation that the accuracy of the estimates decreases as the back-to-back assumption weakens. When the packets are not well correlated, the packets no longer have similar experience on the shared link. In our experiments, with exhaustive trials, the smallest achievable packet spacing is $3\mu s$.

End host load Delays in the generation of SYN-ACK response packets at the end hosts depends on a variety of factors including scheduling and the load. To evaluate these effects, we adjust the end host load by varying the number of active http connections to the Apache 2.0.52 http server on that system. The results of average SYN-ACK generation delay are illustrated in Table I. The variability of the SYN-ACK generation delay is negligible even for large number of active simultaneous connections, *e.g.*, 201 connections as shown in Table I.




Fig. 8. Plot of signal to noise ratio (vertical axis) given different number of packet pair probe measurements.

Number of measurements We compute the SNR for each measurement sets (varying the number of packet pair probe measurements), the SNR increases with the number of probes, as shown in Fig. 8. As expected, the confidence of the estimator increases with increasing number of measurements. Typically, the number of probes needed depends on the variability in traffic and the back-to-back assumption, refer to [14] for computational details. Theoretically, one can use infinite number of measurements to achieve high accuracy. However, in practice, the measurement period increases as we increase the number of probes. The measurement period should always be shorter than the routing table updates. If one increases the probe frequency, it is possible to interfere with the normal traffic and induce congestion on the outbound link. The processing load at the receiver ends might also be increased.

B. Internet Experiments

We randomly chose 6 universities on the Abilene Network and we collect measurements from both Rice University and University of Wisconsin - Madison. The university homepages indexed from 1–6 are Oklahoma University, Virginia Commonwealth University, Vermont University, Utah State University, University of Washington and University of Idaho. We collect a set of estimates with good confidence given by the standard deviation of the estimates. If the ratio between the standard deviation of the estimate and the estimate (SNR) is more than 3 (as recommended in Section IV), we use them as an input to the topology identification algorithm. The inputs to the clustering algorithm are the estimates and its variability, as the mean and the variance. We normalize the variance by the number of measurements. An example of the results of topology measurement experiments based on RTT tomography is shown in Figure 9. The “true” topology is given by traceroute. The estimated topology is very close to what we expect. The estimated topology results as an artifact of a binary tree. These results are representative of our other topology tests (topologies estimated from over 50 homepages) using Network Radar.




Fig. 9. Comparison between (a) “true” topology and (b) estimated topology for 6 universities on the Abilene Network. Probes are sent from Rice University.

VI. CONCLUSIONS

Many of the prior methods for network tomography rely on coordinated measurement infrastructures. This has limited both the share of the Internet over which the tools could be used, and the number of people who could potentially take advantage of this useful inference technique. In this paper we present and evaluate a network tomography method based on the use of round trip time measurements with confidence estimator. The confidence estimator is critical for RTT tomography since probe packets are subject to more dynamic conditions than probes in a one-way coordinated infrastructure. The confidence estimator enables inaccurate delay estimates to be rejected.

We have also detailed some of the operating conditions in which our tool is effective. We summarize them briefly here. The effectiveness of the tool depends on two major components: (1) the ability to place the probe packets closely spaced, (2) correlation on shared link and (3) the unshared link variability, given the fixed number of measurements.

In this paper, we show that the tool, Network Radar, can be conducted anywhere in the Internet the same way as other network diagnostic tools. The tool is ready to be used and it is in the final stage in integrating the ability to compute the confidence interval. The beta version will be available for download in the near future. We have tested the tool in both controlled environment as well as in the wild Internet. We have shown that the confidence interval derived from our estimator provides a good estimate to the traffic condition. By collecting measurements with high confidence interval, one could use it for topology identification.

REFERENCES

- [1] R. Castro, M. Coates, and R. Nowak, “Likelihood based hierarchical clustering,” *IEEE Transactions in Signal Processing*, August 2004.
- [2] Y. Tsang, M. Coates, and R. Nowak, “Network delay tomography,” *IEEE Transactions on Signal Processing*, vol. 51, no. 8, pp. 2125–2136, August 2003.
- [3] V. N. Padmanabhan, L. Qiu, and H. Wang, “Server-based inference of Internet link lossiness,” in *Proceedings of IEEE Infocom*, San Francisco, CA, April 2003.
- [4] M. Coates, A. Hero, R. Nowak, and B. Yu, “Internet tomography,” *IEEE Signal Processing Magazine*, May 2002.
- [5] S. Bohacek and B. Rozovskii, “Models and Techniques for Network Tomography,” in *IEEE Workshop on Information Assurance and Security*, 2001.
- [6] Y. Tsang, M. Coates, and R. Nowak, “Passive network tomography using EM algorithms,” in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Proc.*, May 2001.
- [7] A. Pásztor and D. Veitch, “Precision based precision timing without gps,” in *Proc. ACM SIGMETRICS*, Marina Del Rey, Los Angeles, June 2002.
- [8] D. Veitch, S. Babu, and A. Pásztor, “Robust Synchronization of Software Clocks Across the Internet,” in *Proceedings of ACM/USENIX Internet Measurement Conference*, Taormina, Italy, October 2004.
- [9] Y. Tsang, M. Yildiz, P. Barford, and R. Nowak, “Network radar: Tomography from round trip time measurements,” in *Proceedings of ACM Internet Measurement Conference*, Taormina, Sicily, Italy, Oct 2004.
- [10] K. Anagnostakis, M. Greenwald, and R. Ryger, “cing: Measuring network-internal delays using only existing infrastructure,” in *Proceedings of IEEE Infocom*, San Francisco, CA, April 2003.
- [11] A. Adams, T. Bu, R. Cáceres, N. Duffield, T. Friedman, J. Horowitz, F. Lo Presti, S.B. Moon, V. Paxson, and D. Towsley, “The use of end-to-end multicast measurements for characterizing Internet network behavior,” *IEEE Communications Magazine*, May 2000.
- [12] K. Lai and M. Baker, “Measuring link bandwidths using a deterministic model of packet delay,” in *Proc. ACM SIGCOMM*, Stockholm, Sweden, Aug. 2000.
- [13] N. Spring, D. Wetherall, and T. Anderson, “Scriptroute: A Public Internet Measurement Facility,” in *Proceedings of USENIX Symposium on Internet Technologies and Systems (USITS)*, Seattle, WA, March 2003.
- [14] Y. Tsang, M. Yildiz, P. Barford, and R. Nowak, “Network tomography unleashed,” in *Technical Report*, University of Wisconsin – Madison, 2005.
- [15] P. Barford and L. Landweber, “Bench-style network research in an Internet instance laboratory,” *Computer Communications Review*, vol. 33(3), July 2003.
- [16] J. Sommers and P. Barford, “Self-configuring network traffic generation,” in *Proceedings of ACM Internet Measurement Conference*, Taormina, Sicily, Italy, October 2004.
- [17] E. Kohler, R. Morris, B. Chen, J. Jannotti, and F. Kaashoek, “The click modular router,” *ACM Transactions on Computer Systems*, vol. 18(3), August 2000.
- [18] Endace Measurement Systems, “DAG Network Monitoring Interface Cards,” <http://www.endace.com>, 2004.