One-Way Delay Estimation Using Network-Wide Measurements

Omer Gurewitz, Israel Cidon and Moshe Sidi
Electrical Engineering Department
Technion, Haifa 32000, Israel

***Abstract*— We present a novel approach for the estimation of one-way delays between network nodes without any time synchronization in the network. It is based on conducting multiple and simple one-way measurements among pairs of nodes, and estimating the one-way delays by optimizing the value of a global objective function that is affected by the overall network topology and not just by individual measurements. We examine two objective functions. The first intuitive choice is the least square error (LSE). Using a novel concept of delay induced link probabilities, we develop a second objective function that is based on the maximum entropy (ME) principle. Extensive numerical experiments show that both functions considerably outperform the common method of halving the round trip delays. They also show that ME outperforms the commonly used LSE.**

EXTENDED ABSTRACT

Accurate measurements and adequate analysis of network characteristics are essential for robust network performance and management. Such real-data analysis plays a key role in network design and in the control of its dynamic behavior. One of the most important network performance quantities is delay as it strongly influences the configuration and performance of network protocols such as routing and flow control and network services such as voice and video over IP. Delay measurements are common in such environments and many others. Furthermore, continuous monitoring of delay is essential in many applications in order to check compliance with critical delay constraints.

In many cases the path from a source to a destination may differ from the path from the destination back to the source. Even when the two paths are symmetric, they may have different performance characteristics due to asymmetric loads or different QoS provisioning [1], [2]. Moreover, performance of many applications depends mostly on the delays in one direction [3]. For example, streaming applications performance depends more on the characteristics of the path from the source to destination. A typical client server transaction depends more on the quality of the path from the server to the client. Finally, for voice and video conferencing each unidirectional path is responsible for timely delivery. Consequently, the capability to measure or estimate one-way delays is very important.

The main obstacle in measuring one-way delays is that clocks in a network are not synchronized. Taking one-way measurements is quite simple. A node can send a probe packet with a time stamp on it to its neighbor. When the neighbor receives the packet, it marks its own time stamp over it. The difference between these two time stamps is a *one-way measurement*. Clearly, this one-way measurement equals the corresponding one-way delay only if the clocks of the two nodes are synchronized. Otherwise, the one-way measurement includes the corresponding one-way delay and the clock offset (that is unknown) between the nodes.

Global Positioning Systems (GPS) provide accurate time synchronization between network nodes; unfortunately, GPS are scarce in computer networks. Moreover, an embedded GPS requires continuous reception of multiple satellites which is hard to accomplish indoors or at secured data centers. Network Time Protocol (NTP) is the current standard for synchronizing clocks, with respect to Universal Time-Coordinated (UTC), in the Internet [4], [5]. NTP measures round-trip delays and uses a halving procedure to estimate the clock offsets . A recent offset synchronization method was suggested in [6] for a Pentium based systems as an alternative to GPS synchronization. However, this method calls for a GPS level synchronized NTP server in the (delay-wise) proximity of the measurement endpoints, a requirement that is not practical many times for remote endpoints, branches and homes and cannot be implemented in non-PC based systems. A novel synchronization protocol based on NTP messages that provides better accuracy by optimizing a global cost function is described in [7]. However, all these clock synchronization procedures are working accurately only when the delay is symmetric. Another approach for synchronizing clocks in sensor networks based on the availability of broadcast and low propagation delay among neighboring sensors appears in [8].

Unlike one-way delay, round-trip delay measurements are simple to conduct and they are accurate since the same clock is used while transmitting the packet and upon its return; a common approach used for estimating one-way delay is to measure round-trip delays and halve them. This requires not only that the route between source and destination be the same, but that traffic loads and QoS configurations in both directions also be the same. However, as noted above, often this is not the case.

We present a novel approach for the estimation of one-way delays from one-way measurements that do not require clock synchronization among the nodes of the network. The approach is based on taking one-way measurements between neighboring nodes and pose these measurements as constraints to well defined optimization problems.

We show that the one-way measurements impose constraints on the feasible values of the one-way delays. Our goal is to derive the 'best' estimate of the one-way delays given the

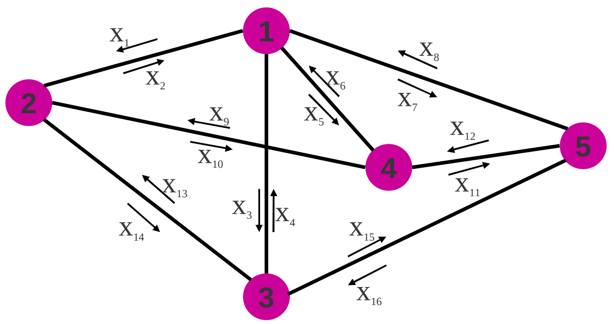

Fig. 1. 5-Node 16-link network

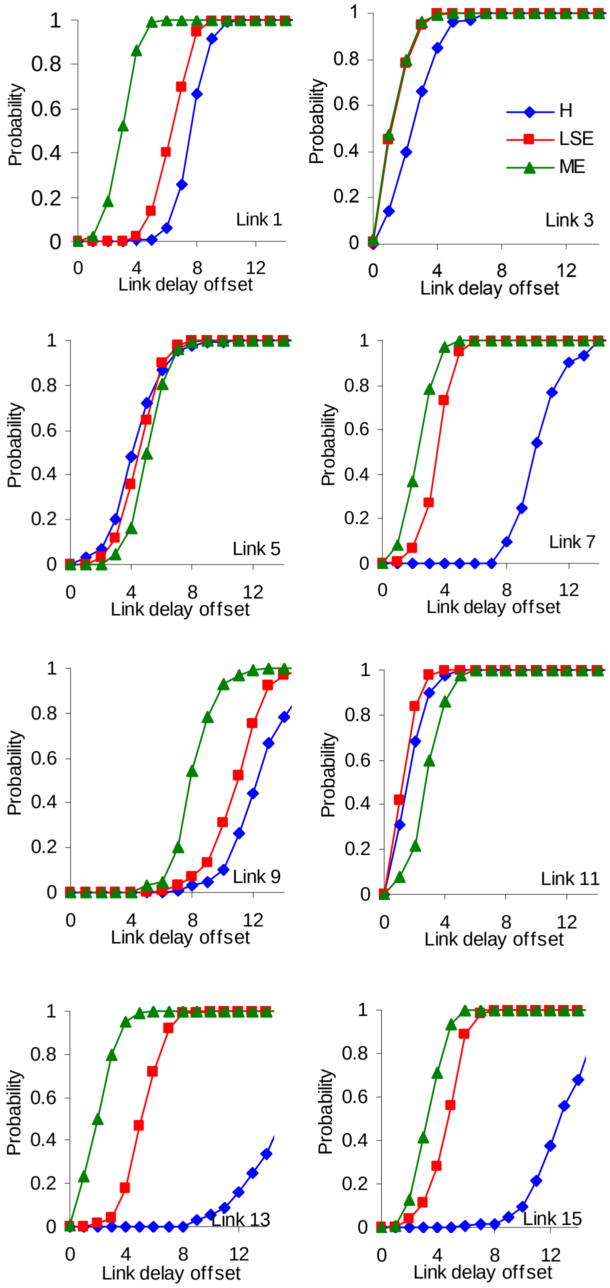

Fig. 2. The fraction of runs that the difference between the estimated link delay and the real link delay is not greater than t, for selected links in the 5-node 16-links network

one-way measurements. We show how to exploit the one-way measurements to obtain the necessary constraints on the one-way delays and derive the number of independent constraints that can be obtained. These constraints are used in the optimization problems that we explore. For these problems we define objective functions that when optimized, provide the 'best' estimate for the one-way delays.

We investigate two different objective functions. The first objective function is very intuitive and is based on the *least square error* (LSE) principle. According to this principle the solution that is sought for is the one that minimizes the square error. The second objective function is based on the *maximum entropy* (ME) principle. According to this principle the solution that is sought for is the one that maximizes the entropy. Note that the definition of entropy requires an underlying probability space. One of our contributions is the introduction of a method to induce probabilities upon network links that are relative to the delays over these links. The objective function based on the ME principle yields itself to relatively simple computations and results in a better one-way delay estimation for most cases checked.

Both objective functions provide estimates of the fixed part (i.e., propagation) of the one-way delay. For the estimation of the variable delay one can use the same optimization methodology and common techniques that are available for the estimation of the distribution parameters. The solutions that are provided are easy to implement using standard probe packets among nodes (e.g., NTP, ICMP). Extensive numerical experiments demonstrate that both schemes considerably outperform the traditional round trip delay halving.

As an example, Figure 2 shows the results of the 100 runs over a variety of selected links of the network in Figure 1. The y axis in each graph presents the fraction of runs where the propagation delay difference between the estimated value and the real propagation delay is not greater than the value described in the x axis. Figure 2 demonstrates significant improvement in terms of the delay estimation of both the "LSE" and "ME" schemes over the halving scheme. For example, in link 7, the estimated link propagation delay never exceed 5.1 and 3.7 time units in 100 runs for the LSE and ME, respectively, whereas for the halving scheme the maximum delay error is 13.8. Symmetric or nearly symmetric delay mean and variance such as in links 5,6 and 11,12 make the halving scheme the natural choice for estimating the one-way link delay. It is important to note that the delay estimation on such links by LSE and ME does not fall much behind the halving scheme. On the other hand, on asymmetric links, the LSE and ME schemes significantly outperforms the halving scheme.

We ran the same simulation as before, 100 times over the 5-node 16-link network where the delay at each link has Normal distribution. We estimated the distribution based on the Mean and Variance. We compared the results of the LSE and ME with the halving technique.

Figure 3 shows that the estimates obtained by the LSE and ME are much better than those obtained by the halving scheme. We added to each graph in this figure the I-divergence

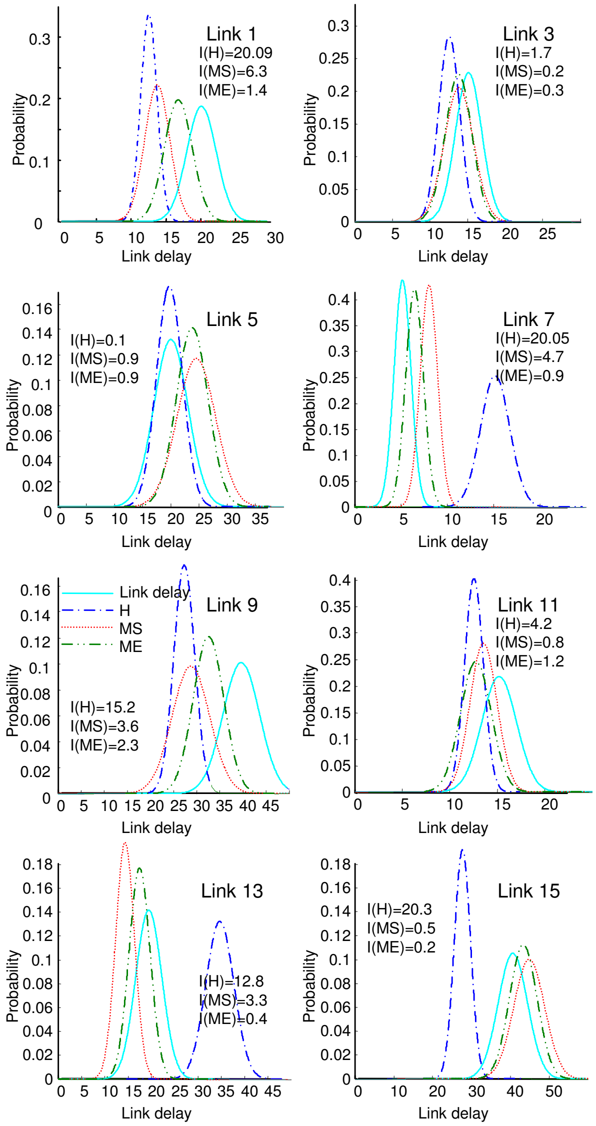

Fig. 3. Estimation of the Normal distribution propagation delay in the 5-node 16-link network. H-Halving, MS-Least Square Error, ME-Maximum Entropy.

distance [9], which measures the difference between the estimated and true distribution. It is interesting to note the results of link 5, which was forced to be symmetric, i.e., to have the same Mean and Variance as link 6. The halving scheme is naturally the best for symmetric links. However, the difference from the two schemes is not very big. For link 5 the I-divergence distance is 0.1 and 0.9 from the link delay to the estimation based on halving and both LSE and ME, respectively. On the other hand, in links that are not symmetric, the improvement of LSE and ME over the halving scheme is very significant. For instance, for link 1 the I-divergence distance is 20.9, 6.3 and 1.4 from the link delay halving, LSE and ME, respectively. Again, ME outperforms the LSE scheme.

ACKNOWLEDGMENTS

The authors would like to thank Prof. Hanoch Levy for helpful discussions.

This research has been made possible in part by a grant from the Cisco University Research Program Fund at Community Foundation Silicon Valley.

REFERENCES

[1] V. Paxson. End-to-end internet packet dynamics. *IEEE/ACM Transactions on Networking*, 7(3):277–292, 1999.

[2] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, Deb Moll, R. Rockell, T. Seely, and C. Diot. Packet-level traffic measurements from the Sprint IP backbone. *IEEE Network*, 17(6):6–16, 2003.

[3] G. Almes, S. Kalidindi, and M. Zekauskas. A one-way delay metric for IPPM. RFC 2679, September 1999.

[4] D.L. Mills. Improved algorithms for synchronizing computer network clocks. *IEEE*.

[5] D.L. Mills. Network Time Protocol (version 3) specification, implementation and analysis. Network working group report RFC-1305, University of Delaware, 1992.

[6] D. Veitch, S. Babu, and A. Pasztor. Robust synchronization of software clocks across the internet. In *Internet Measurement Conference*, Taormina, Italy, October 2004.

[7] O. Gurewitz, I. Cidon, and M. Sidi. Network time synchronization using clock offset optimization. In *ICNP 2003*, pages 212–221, Atlanta, GA, November 2003.

[8] J. Elson, R. Karp, C. Papadimitriou, and S. Shenker. Global synchronization in sensornets. In *Proc. of the 6th Latin American Symposium on Theoretical Informatics (LATIN'04)*, pages 609–624, Buenos Aires, Argentina, April 2004.

[9] D. Kazakos and P. Papantoni-Kazakos. *Detection and estimation*. Computer Science Press, Rockville, MD, 1990.