

Secure Latency Estimation with Treeples

Eric Chan-Tin and Nicholas Hopper
University of Minnesota
Minneapolis, MN
USA

{dchantin, hopper}@cs.umn.edu

ABSTRACT

A network latency estimation scheme associates a “position” to every peer in a distributed network such that the latency between any two nodes can be accurately estimated from their positions. Applications for these schemes include efficient overlay construction, compact routing, anonymous route selection, and efficient byzantine agreement. We present a new latency estimation scheme, Treeples. Our scheme is different from existing ones in several aspects: Treeples is provably secure, rather than being able to resist known attacks; positions in Treeples are not Euclidean coordinates and reflect the underlying network topology; finally, positions in Treeples are accurate, stable, and can be assigned to peers not participating in the system.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General—*Security and protection*

General Terms

Security, Theory

Keywords

network latency estimation, secure

1. INTRODUCTION

Each node in a network coordinate system [5, 7, 17–19, 22] gets assigned a *coordinate* such that the distance between any two nodes’ coordinates is a good estimate of the Internet round-trip time between those two nodes. The accurate estimation of network distance between arbitrary nodes is useful in many situations: finding the closest node to download content from in a content distribution network or file-sharing system [25]; reducing inter-ISP communication [4, 14]; reducing state in internet routers [1, 9, 13]; detecting Sybil attacks [2, 8]; and conducting byzantine leader elections [6]. Early coordinate systems have been shown to exhibit good accuracy and fast convergence to stable coordinates.

However, Kaafar *et al.* [10, 11] showed that these early systems were vulnerable to simple attacks, where the adversarial nodes can disrupt the system by claiming to have randomly-chosen coordinates, and adding random delays to

their outgoing messages. To mitigate these attacks, several schemes have been proposed [12, 21, 23, 26, 27]. However, none of them have any explicit security goal, and at least one of these schemes [27] has been shown to be insecure against more subtle adversaries [3]. This can lead to a potential “penetrate and patch” cycle of designs and attacks.

In this proposal, we introduce a strong definition of security for network latency estimation systems which is robust to new attacks. An adversary should not be able to influence the estimated distance between two honest nodes, and should only be able to inflate the distance between itself and another honest node. We also show that schemes using trusted landmark nodes [17] can also fail by allowing malicious peers to reduce their distance to targeted victim nodes. The common fault from this vulnerability and the frog-boiling attack [3] is that current network coordinate systems treat the underlying network topology as a black box – this makes it hard to differentiate anomalous RTTs, which can be caused by changing network conditions, or due to adversarial manipulation. We show that our network latency estimation scheme, Treeples, is secure under our security goals, and is also accurate. Using a large real-world dataset [16], Treeples has a median relative error of 0.26, while Vivaldi [7], a popular insecure network coordinate system, has a median relative error of 0.25. This means that 50% of the estimated distances are within 26% of the real network latency. Moreover, Treeples network positions are highly stable – positions calculated on day one can be used with the same accuracy 21 days later. Thus, peers do not need to constantly compute their positions, reducing bandwidth consumption.

2. SECURITY

2.1 Threat Model

We assume the network consists of a collection of *end-hosts* connected by *routers*. An adversary can control an arbitrary number of end-hosts, but not routers. Thus an adversary can communicate with any host, fake the origin of messages, deviate from protocols, collude, and inflate the measurement of $r_{tt}(n, m)$ where m is malicious. However, the adversary cannot affect the r_{tt} between honest hosts. Although it may seem that excluding routers from being malicious is a strong assumption, we argue that excluding routers is reasonable in our model: if an adversary could arbitrarily delay or redirect messages, then a network latency estimation scheme cannot succeed, since *any* latency estimate can be invalidated by the adversary. While the most desirable situation would be to allow routers to become adversarial, all the currently known

attacks fit within our threat model. Therefore, a scheme that is provably secure under our threat model is secure against all the currently known attacks.

2.2 Definitions

Accuracy: A scheme is accurate if the computed estimated distance between two nodes is close to the real network distance between these two nodes. The *median relative error* is commonly used in the literature to measure accuracy. The median relative error of peer P is

$$\text{median}_{P' \in \text{Peers}} \frac{|\text{Distance}(\rho(P), \rho(P')) - \text{rtt}(P, P')|}{\text{rtt}(P, P')},$$

where $\rho(P)$ is the coordinate of peer P . If a network coordinate system has a median relative error c , then this implies that 50% of the estimated distances are within c of the real network distance.

Security Goal: Our two security goals can be described as follows: 1) An adversary cannot affect the distance between two honest nodes, and 2) An adversary can only inflate (increase) the distance between itself and another node.

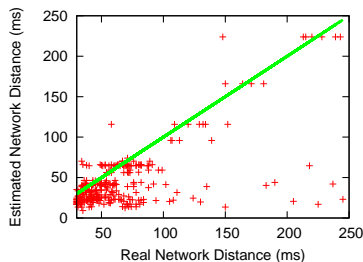


Figure 1: Targeted close node attack results on GNP

2.3 Failures of Previous Coordinate Systems

Secure decentralized network coordinate systems have been shown to be vulnerable against the frog-boiling attack [3]. Centralized network coordinate schemes consist of trusted nodes which compute the coordinates of other nodes. GNP [17] is such a scheme. Although the authors do not claim their scheme is secure against attacks, it seems that intuitively if nodes obtain digitally signed messages from each landmark attesting to RTT measurements, GNP could be secure. However, we show that it is possible for an adversary to influence the system so as to artificially decrease the estimated distance between itself and a targeted honest node. In our simple attack, the malicious node a knows the coordinate of the target node t , so a can compute the rtt between t and each landmark node. When a landmark ℓ_i measures $\text{rtt}(\ell_i, a)$, the adversarial node a attempts to make the measurement as close to $\text{rtt}(\ell_i, t)$ as possible, subject to the constraint that the RTT cannot be smaller than the real network distance. We repeated this experiment using the code and matrix topology from [24]. For each simulation run, we randomly pick one target node and one malicious node from a set of 101 nodes, where 10 of them were considered to be landmarks. Figure 1 shows that our attack works: in nearly all the runs, a receives a coordinate that is closer to t than the real network distance.

3. TREEPLE

To illustrate the idea for Treeple, suppose we have a system with trusted vantage point A and peers X , Y , and Z . A




Figure 2: A tree built from one source to three destinations with link latency in milliseconds.

can find the network paths (using traceroute or other similar technique) from itself to each of the three peers, constructing a tree with routers as the internal nodes and link latencies for the edges. The vantage point A can then use this tree to calculate an upper bound on the estimated network latency between X and Y : first find C , the least common ancestor of X and Y ($\text{lca}(X, Y)$). Then calculate the sum of the link latencies between C and X and between C and Y . Due to the Internet routing policies, the RTTs measured might not necessarily reflect the latencies exactly due to asymmetric return paths; also the *triangle inequality violation* [15,28] might occur, and the path $X-C-Y$ will be shorter than the actual network path. However, we show that our scheme still provides a good estimate for network latency. For example, let's assume A builds the tree as shown in Figure 2. The $\text{lca}(Y, Z) = B$ and the estimated distance between Y and Z is $35 + 15 + 20 = 70$ ms. In this scheme, a malicious peer X can only affect the distance (real or estimated) between itself and another peer (Y or Z), but it cannot affect the distance between Y and Z since it is not on the same network path.

Using only one vantage point is most likely not desirable as the tree distance for some end-hosts might be larger than the real network distance, since many network links will not be included with one vantage point. To address this issue, we choose k topologically distinct vantage points. An end-host coordinate then becomes the route from each trusted vantage point, and the distance between two end-hosts is the minimum of the k distances.

THEOREM 1. *Assuming the set $\{T_1, \dots, T_k\}$ of vantage points are honest, Treeple is a secure network latency estimation scheme.*

It is intuitively easy to see that Treeple is secure under our security goals. We omit the proof due to space constraints.

4. EVALUATION

Experimental setup: To evaluate our proposed scheme, we used the iPlane [16] dataset. It contains the results of periodic traceroutes from about 250 PlanetLab [20] nodes to thousands of other IP addresses. We used the datasets from Dec 1st to Dec 22nd, 2009.

Selecting vantage points: We chose different sizes k of vantage points using a *greedy sampling* algorithm, which works as follows. First, we selected at random a set S of 1,000 pairs (A, B) for which we had latency measurements. We then picked the best vantage point T_1 for the pairs in S

among the 250 possible choices. We then picked the second best vantage point T_2 such that combined with T_1 would produce the lowest median relative error for S , and so on until k vantage points have been chosen.

Baseline: Vivaldi. Vivaldi [7] is a popular network coordinate system. The Vuze file-sharing application [25] implements a similar network coordinate scheme and Vivaldi is used for computing coordinates in various “secure” coordinate schemes [12, 23, 27]. Using the iPlane dataset, we ran a simulation of Vivaldi. The median relative error obtained was 25%.




Figure 3: CDF of relative error when varying the number of vantage points k used

Accuracy: Figure 3 shows the CDF of the relative error for varying number of vantage points k . We note that using fewer than 5 vantage points produces the same CDF curve. Increasing k from 5 to 10 increases the accuracy as the median relative error is decreased. Increasing k past 10 produces the same median relative error but the gain in accuracy can be seen at the 90th percentile. Increasing the number of vantage points increases the accuracy of our scheme – the median relative error is 26% (compared to 25% for Vivaldi). From the figure, only 20 trusted vantage points are needed to be able to accurately estimate the network distance between two nodes.

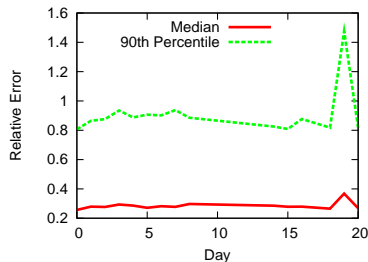


Figure 4: Median and 90th percentile relative error using 12/01/09 positions for estimation measurements through 12/21/09, by day.

Stability: If end-hosts in Treeples need to contact the vantage points often to obtain their new positions, then these nodes become a central point of failure as they would have to be online all the time. We now consider whether the accuracy of the Treeples’s positions changes over time. We used the best 20 trusted nodes from 12/01/2009 to estimate the network distances for the nodes from 12/01/2009 to 12/21/2009 (three weeks). The median and 90th percentile relative errors are shown in Figure 4. This shows that Treeples remains accurate over time. Thus, frequent network measurements are not needed: the same positions can be used weeks later.

5. FUTURE WORK

There are possible improvements to Treeples which could further improve accuracy. It is known that multiple interfaces, load balancers, multi-protocol label switching, and hosts which do not respond, can affect the accuracy of network paths returned by traceroute. Resolving them could improve the accuracy of Treeples. Moreover, we treated the trusted vantage points as fixed. However, it is possible to have the vantage points migrate using standard methods such as migrating trusted servers for services for DNS, DHTs, and Tor. Finally, we plan to consider router compromise and expand our threat model to determine if Treeples is still provably secure if, for example, BGPsec is deployed.

6. REFERENCES

- [1] Ittai Abraham and Dahlia Malkhi. Compact routing on euclidian metrics. In *PODC '04: Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*, pages 141–149, New York, NY, USA, 2004. ACM.
- [2] R. Bazzi and G. Konjevod. On the Establishment of Distinct Identities in Overlay Networks. In *ACM PODC*, 2005.
- [3] Eric Chan-Tin, Daniel Feldman, Yongdae Kim, and Nicholas Hopper. The Frog-Boiling Attack: Limitations of Anomaly Detection for Secure Network Coordinates. *SecureComm*, 2009.
- [4] D. Choffnes and F. Bustamante. Taming the Torrent: A practical approach to reducing cross-ISP traffic in P2P systems. *ACM Special Interest Group on Data Communication (SIGCOMM)*, 2008.
- [5] M. Costa, M. Castro, A. Rowstron, and P. Key. PIC: Practical Internet Coordinates for Distance Estimation. *Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2004.
- [6] James Cowling, Dan Ports, Barbara Liskov, Raluca Ada Popa, and Abhijeet Gaikwad. Census: Location-Aware Membership Management for Large-Scale Distributed Systems. In *the proceedings of USENIX Annual Technical Conference*, 2009.
- [7] Frank Dabek, Russ Cox, Frans Kaashoek, and Robert Morris. Vivaldi: A Decentralized Network Coordinate System. In *Proceedings of ACM SIGCOMM*, 2004.
- [8] John R. Douceur. The sybil attack. In *Proc. of the International Workshop on Peer-to-Peer Systems (IPTPS)*, 2002.
- [9] R. Gummadi, R. Govindan, N. Kothari, B. Karp, Y. J. Kim, and S. Shenker. Reduced state routing in the internet. *HotNets*, 2004.
- [10] M. A. Kaafar, L. Mathy, T. Turlletti, and W. Dabbous. Real attacks on virtual networks: Vivaldi out of tune. *Proceedings of the SIGCOMM workshop on Large-scale Attack Defense*, 2006.
- [11] M.A. Kaafar, L. Mathy, T. Turlletti, and W. Dabbous. Virtual Networks under Attack: Disrupting Internet Coordinate Systems. *ACM/e-NEXT International Conference on Future Networking Technologies (CoNext)*, 2006.
- [12] Mohamed Ali Kaafar, Laurent Mathy, Chadi Barakat, Kave Salamatian, Thierry Turlletti, and Walid Dabbous. Securing Internet Coordinate Embedding Systems. *Proceedings of ACM SIGCOMM*, 2007.
- [13] Jonathan Ledlie, Michael Mitzenmacher, and Margo Seltzer. Wired geometric routing. *International Workshop on Peer-to-Peer Systems (IPTPS)*, 2007.
- [14] C. Lumezanu, D. Levin, and N. Spring. Peer wise discovery and negotiation of faster path. In *Proceedings of HotNets-VI*, 2007.
- [15] Cristian Lumezanu, Randy Baden, Neil Spring, and Bobby Bhattacharjee. Triangle inequality variations in the internet. In *IMC '09: Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, pages 177–183, New York, NY, USA, 2009. ACM.
- [16] Harsha V. Madhyastha, Tomas Isdal, Michael Piatek, Colin Dixon, Thomas Anderson, Arvind Krishnamurthy, and Arun Venkataramani. iPlane: An Information Plane for Distributed Services. In *Proceedings of the USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2006.
- [17] T. S. Eugene Ng and Hui Zhang. Predicting Internet Network Distance with Coordinates-Based Approaches. *Proceedings of IEEE INFOCOM*, 2002.
- [18] T. S. Eugene Ng and Hui Zhang. A network positioning system for the internet. *Proceedings of the USENIX annual technical conference*, 2004.
- [19] M. Pias, J. Crowcroft, S. Wilbur, T. Harris, and S. Bhatti. Lighthouses for Scalable Distributed Location. *International Workshop on Peer-to-Peer Systems (IPTPS)*, 2003.
- [20] PlanetLab. <http://planet-lab.org>.
- [21] Damien Saucez, Benoit Donnet, and Olivier Bonaventure. A Reputation-Based Approach for Securing Vivaldi Embedding System. In *EUNICE Open European Summer School and IFIP TC6.6 Workshop on Dependable and Adaptable Networks and Service*, 2007.
- [22] Y. Shavitt and T. Tankel. Big-Bang Simulation for embedding network distances in Euclidean space. *IEEE INFOCOM*, 2003.
- [23] Micah Sherr, Matt Blaze, and Boon Thau Loo. Veracity: Practical Secure Network Coordinates via Vote-based Agreements. In *USENIX Annual Technical Conference*, 2009.
- [24] GNP Simulator. <http://www.cs.rice.edu/~gw4314/ncs-configurable.tar.gz>.
- [25] Vuze. <http://azureus.sourceforge.net>.
- [26] Guohui Wang and T. S. Eugene Ng. Distributed Algorithms for Stable and Secure Network Coordinates. *ACM/USENIX Internet Measurement Conference (IMC)*, 2008.
- [27] David Zage and Cristina Nita-Rotaru. On the accuracy of decentralized virtual coordinate systems in adversarial networks. In *Proceedings of the 14th ACM conference on Computer and communications security (CCS)*, 2007.
- [28] Han Zheng, Eng Keong Lua, Marcelo Pias, and Timothy Griffin. Internet Routing Policies and Round-trip Times. *Passive and Active Measurement Workshop*, 2005.