Impact of Overlay Routing on End-to-End Delay

Hui Zhang[1,2], Li Tang[1,2], Jun Li[2,3]
[1] Department of Automation, Tsinghua University, Beijing, China
[2] Research Institute of Information Technology, Tsinghua University, Beijing, China
[3] Tsinghua National Lab for Information Science and Technology, Beijing, China
zhanghui04@mails.tsinghua.edu.cn

Abstract—**The prosperity of various real-time applications triggers significant challenges for the Internet to meet the critical end-to-end delay requirements. This paper investigates the feasibility and practical issues of using overlay routing to improve end-to-end delay performance, with the analysis of round trip time data collected over three months by the All Pairs Pings project between each pair of hundreds of nodes on the PlanetLab. The results show: 1) overlay routing has the potential to reduce the round trip time by 40 milliseconds and increase network connectivity by 7% on average; 2) even simple static overlay paths can reduce end-to-end delay, while the dynamic algorithm does better; 3) Over 80% of the shortest overlay paths have no more than 4 hops, and a simple algorithm leveraging only one relay node can efficiently take the advantages of overlay routing.**

Keywords—Overlay Routing, End-to-End Delay

I INTRODUCTION

Nowadays more and more applications on the Internet have triggered critical delay challenges. For example, voice over IP (VoIP) requires one-way end-to-end (E2E) delay under 150ms [1] and some real-time multiplayer online games such as Quake III require delay under 100ms [2] to achieve favorable user experience. Compared with traditional applications such as FTP and HTTP, some of these interactive applications require low bandwidth and tolerate high packet loss relatively. For example, one ITU-T G.723.1 voice stream only takes up bandwidth of less than 10Kbps, and G.711 voice stream can keep high quality under 5% random packet loss [3]. On the other hand, however, E2E delay does not benefit from the improvement of bandwidth that is doubling every 1.9 years [4], but in contrast becomes even worse along with the Internet's expansion. Thus it is more urgent to improve the E2E delay performance than bandwidth and packet loss to ensure the quality of service (QoS) of such delay-critical applications.

The Internet provides a successful set of IP-layer protocols contributing to its global deployment, reliability and standardization. With the unprecedented growth of Internet, these protocols can hardly afford major changes. As it is extremely difficult to implement new mechanisms at IP-layer, overlay routing offers a way to improve the performance of delay-critical application without the need to modify the existing protocols and infrastructure. Prior work suggested that routing overlays can enable the Internet to achieve a higher degree of application performance and reliability than current IP-layer only routing machanisms. The objective of this paper is to investigate whether and to what extent overlay routing can improve E2E delay. To achieve this objective, we analyze the data of round trip time (RTT)[1] collected over three months by the All Pairs Pings (APP) project [5] between each pair of hundreds of nodes on the PlanetLab [6]. The results not only affirm the effectiveness of overlay routing to improve E2E delay, but also indicate some valuable principles for the design of overlay routing algorithms.

The rest of this paper is organized as follows: Section II introduces the dataset and methodology. Section III gives an overview of using overlay routing to improve E2E delay. Section IV deeply investigates the overlay routing performance and brings forward some principles for the design of practical overlay routing algorithms. Section V reviews related work. Section VI concludes the paper and presents some future work.

II DATASET

The PlanetLab was chosen as the testbed for our investigation, because it represents a typical Internet environment with hundreds of nodes distributed worldwide. The PlanetLab APP project collects minimum, average, and maximum pairwise RTT by 10 consecutive ping results, repeating approximately every 15 minutes. Measurements are taken locally from individual production nodes perspective, stored locally, and periodically archived centrally at MIT. If no pings have successfully returned within two minutes, the link is considered to be down, and a link failure is recorded correspondingly.

The dataset in the paper is collected from June 1, 2005 to August 31, 2005, including 8813 archives, 478 nodes in each archive on average and nearly 1,400,000,000 RTT tuples consisting of the minimum, mean, and maximum. The following investigation is based on the RTT matrix formed by the mean RTTs extracted from each archive.

---

[1] E2E delay means RTT in the following discussion.

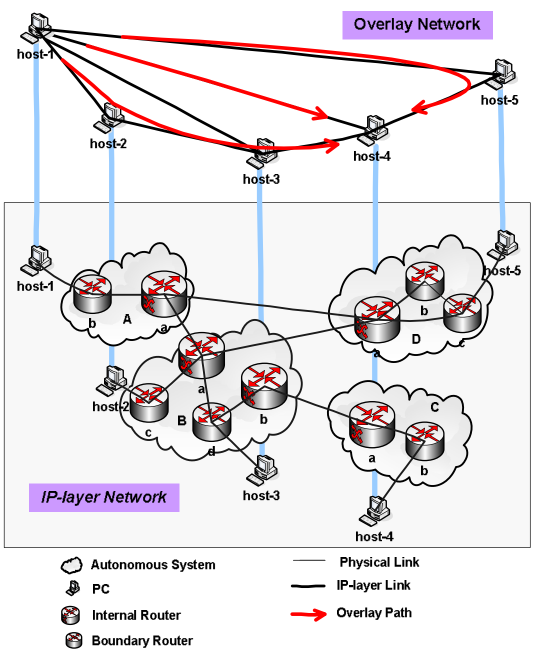

Fig. 1 Overlay Network and IP-layer Network

Theoretically, the E2E delay behavior between a pair of nodes in the Internet is a continuous process, and the data shared by the APP project can be taken as discrete samplings of the process with an interval of 15 minutes. Distinctively, each of the sampling result here is not a single RTT value but a statistical tuple of no more than 10 consecutive measurements. In this paper, we use the mean RTT value in each interval as the sampling value. This approach makes little difference to the inherent nature of E2E delay behavior, because the RTT measurements in an interval are rather concentrated.

Given the minimum, mean, and maximum of a batch of RTT measurements in an interval, the upper bound of their standard deviation can be calculated out according to the Theorem I (see [7] for proof).

Theorem I: If the minimum, mean, and maximum of arbitrary positive real numbers $x_1, x_2, \cdots, x_n$ are respectively $\underline{x}$, x and $\overline{x}$, then the upper bound of their standard deviation is

$$u = \sqrt{\frac{1}{n-1} \cdot \left[p(\underline{x} - x)^2 + q(\overline{x} - x)^2 + (n-p-q)(y-x)^2 \right]},$$

where $p = \left\lfloor n \cdot \frac{\overline{x} - x}{\overline{x} - \underline{x}} \right\rfloor, q = \left\lfloor n \cdot \frac{x - \underline{x}}{\overline{x} - \underline{x}} \right\rfloor, y = n \cdot x - p \cdot \underline{x} - q \cdot \overline{x}$.

When using α, the ratio of standard deviation to mean, as the divergence criteria for a specific batch of RTT measurements, we find that over 80% of RTT tuples in the data have their α less than 0.1, and over 90% less than 0.2. Therefore, the mean of the RTT measurements in an interval is proved to characterize closely the sampling of the E2E delay process at that time.

III PERFORMANCE OVERVIEW

In this section, we investigate whether and to what extent an overlay network could improve E2E delay on average, providing ideal path selection, mainly from two aspects—delay reducing and connectivity increasing.

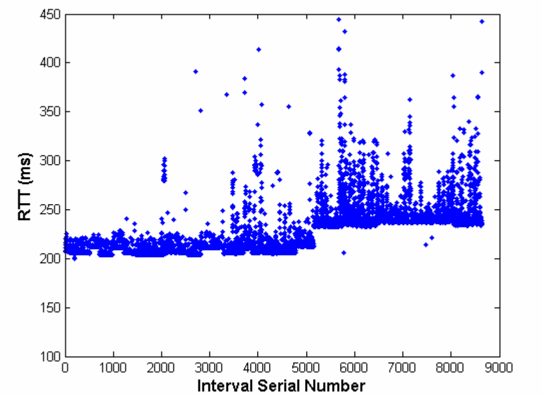

Fig. 2 Time series plot of RTTs between a pair of PlanetLab nodes

A Overlay Network Model

An overlay network constructs an application layer graph on top of an existing IP-layer network, using only a subset of the available network links and nodes, as shown in Fig. 1. An overlay path is a virtual edge in the connectivity graph and may consist of many IP-layer links in the underlying network. Overlay nodes act as application layer routers, forwarding packets to the next hop IP-layer link toward the destination. At the IP-layer, packets traveling along physical links follow the actual IP routing protocol to form each IP-layer link accordingly.

The overlay topology of the PlanetLab in an interval can be modeled as a Weighting Directed Graph (WDG) $G(A_1, A_2, \cdots, A_n)$, where vertex A_i represents a production node in the IP list of an archive, and the weighting of directed edge $\overrightarrow{A_i A_j}$ represents the IP-layer RTT from source node A_i to destination node A_j in the RTT matrix, where $1 \le i, j \le n$. The directed edges with missing data in the RTT matrix are considered to be disconnected or weighting infinite.

Therefore, overlay routing problem can be specific as finding one or more paths through some relay nodes between any pairwise vertexes with low sum of weighting from a series of dynamic WDGs. The found path is defined as overlay path, and the sum of the weightings of directed edges is defined as overlay RTT. In fact, IP-layer path is a special overlay path with no relay node.

B Shortest Path Routing

For a given WDG, the shortest overlay path between any pair of nodes could be calculated out by Floyed-Warshall Algorithm [8, 9], under which network connectivity achieves its upper bound and the E2E delay between any pairwise nodes achieves its lower bound. Hence shortest path routing algorithm is also called ideal overlay routing algorithm.

From the measurement point of view, the E2E delay over a pair of nodes is best separated into two components, a deterministic component and a stochastic component. The stochastic component such as delay introduced by routing alteration, link load and processor load fluctuation makes the E2E delay over a fixed pair of nodes varies significantly with time (see Fig. 2). Whereas, not all the RTTs can accurately reflect the property feature of the path, especially the ones distribute

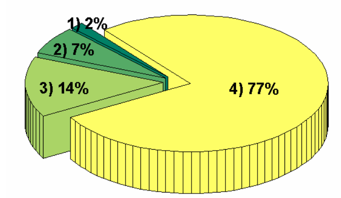

Fig. 3 Proportion of four Categories Pairwise Nodes

at the "tail"[2], for they are likely to be accidental perturbation and only last for seconds. But these "fly points" may produce over optimistic estimation for the performance improvement by overlay routing rather heavily. So the average value of a RTT time series between a pair of nodes in the three months is introduced to represent the weighting of relevant directed edge, aiming at eliminating churn factors. The obtained graph is defined as Average Weighting Directed Graph (Average WDG), and the graph obtained from original RTT matrix is defined as Real-time Weighting Directed Graph (Real-time WDG).

After running Floyed-Warshall Algorithm to the Average WDG, all the pairwise nodes can be classified into the following four categories according to the contrast between IP-layer RTT and overlay RTT, as shown in Fig.3:

1) Both IP-layer RTT and overlay RTT are infinite.
2) IP-layer RTT is infinite, whereas overlay RTT is finite. The study in [11] calls this phenomena to be non-transitivity. Its author explains that, on the PlanetLab, some Internet2-only nodes are not able to directly communicate with Internet1-only nodes and vice versa, while there exists some university nodes that are able to communicate on both Internet1 and Internet2, leading to non-transitivity.
3) IP-layer RTT is equal to overlay RTT.
4) Overlay RTT is lower than the IP-layer RTT. This anti-triangle inequation phenomena rises mainly from the cases of pathological routing between ISPs in current BGP based routing. The MIT CSAIL report [12] also found the similar phenomena that many paths originated in Gigamedia, Taipei and end in other locations in Asia are routed through California in IP-layer, and there usually exists an overlay path via China Telecom, Shanghai, which is more than 50% faster.

As shown in Fig. 3, shortest path rouitng could reduce the pairwise nodes unable to communicate from 9% to 2%, and improve E2E delay for as many as 77% pair of nodes. The average E2E RTT can reduce from 220ms to 181ms. To view these improvements more elaborately, we group the IP-layer RTT by every 50ms and calculate overlay RTT improvement for each IP-layer RTT. As shown in **Fig. 4**, there's hardly any improvement for pairewise nodes with the IP-layer RTT less than 50ms, but the improvement grow rapidly with the increasing of the IP-layer RTT, over 40ms for the IP-layer RTT above 300ms.

IV IN-DEPTH EVALUATION

In this section, we thoroughly investigate overlay routing from the following aspects: 1) How does a static overlay routing perform and is it essential to bring forward dynamic overlay routing? 2) What is the cost of shortest path routing and are there some strategies to reduce the cost? 3) To what extent can multiple paths help to improve the performance of overlay routing?

[2] According to the work in [10] , 84% of E2E delay process a Gamma-like shape with heavy tail.

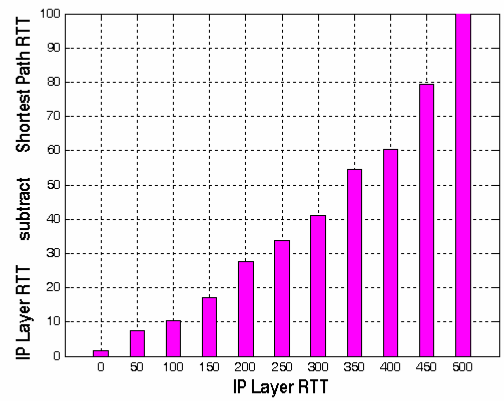

Fig. 4 Delay Improvement by Shortest Path Routing

A Static Overlay Routing and Dynamic Overlay Routing

It should be pointed out that the analysis in section III only shows the performance of a fixed shortest overlay path in the sense of statistics. This subsection focuses on the effects of network dynamics to the overlay routing by comparing the static overlay paths getting from the above Average WDG and the dynamic shortest paths getting from Real-time WDG. The relay nodes on the former paths are obtained form Average WDG by Floyed-Warshall Algorithm mentioned in section III and changelessly during the three month, but the IP-layer RTT of each directed edge varies actually with time. This scheme is defined as Static Shortest Path Routing Algorithm (Static SPRA). Comparatively, the latter paths are obtained from Real-time WDG by Floyed-Warshall Algorithm and defined as Dynamic Shortest Path Routing Algorithm (Dynamic SPRA).

The Cumulative Distribution Function (CDF) curve of overlay RTT to IP-layer RTT in Fig.5 (a) shows, in nearly 70% cases, static shortest path is shorter than IP-layer path; in 20% cases, static shortest path approximates IP-layer path; but in the remaining 10% cases, static shortest paths are worse than IP-layer path. However, take the improvement degree into consideration, Dynamic SPRA is much superior to Static SPRA. For example, about 30% of IP-layer RTTs could be reduced 20% by Dynamic SPRA, but only 10% could be reduced 20% by Static SPRA. Fig.5 (b) also shows Dynamic SPRA can save over 20ms RTT than Static SPRA for over 20% of pairwise nods.

In the aspect of network connectivity, Dynamic SPRA could increase it by nearly 10% in almost every interval, which is 3% higher than the above 7% result in section III. This proves that besides overcoming non-transitivity, overlay routing also helps to detect and recover from a link failure. However, Static SPRA will not improve the network connectivity but serious decrease it by as much as 25%, as seen in Fig. 6, because the static shortest path is not always available during the three months. Once a IP-layer link or a relay node along the static shortest path is failed, connection from source node to destination node will be down. The more hops the static shortest path has, the higher probability the path will break off.

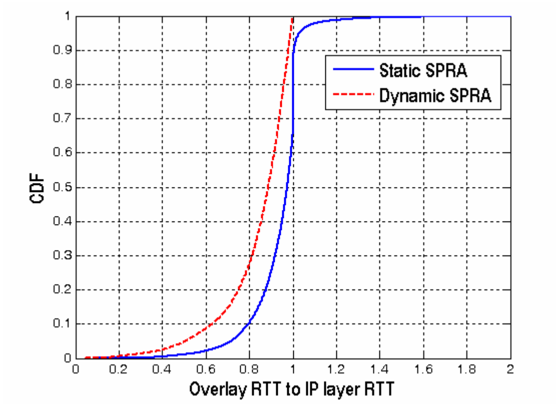

(a) CDF Curve of Relative RTT (Overlay RTT to IP-layer RTT)

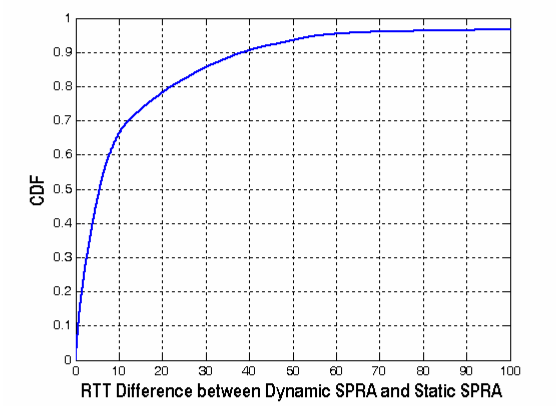

(b) CDF Curve of RTT Difference between Dynamic and Static SPRA

Fig. 5 Comparison between Static SPRA and Dynamic SPRA

From above discussion, it can be concluded that though Static SPRA is effective for some pairs of nodes, the rather poor reliability makes it impractical. Dynamic SPRA is much superior to Static SPRA in both performance and reliability, but suffers from poor scalability, which will be discussed in next section.

B 2-Hop Overlay Routing

Generally speaking, most of the IP-layer shortest path routing algorithm can be classified into two categories: Distance Vector Algorithm and Link State Algorithms. Distance Vector Algorithm requires each router send its entire routing table to its directly connected neighbors on a periodic basis, such as EIGRP, RIP and IGRP. Link State Algorithm floods trigger updates throughout the network only when a network change has occurred, such as OSPF and IS-IS. Assume there are n nodes in an existing network, and the average neighbor number for each nodes is k. Both of Distance Vector Algorithm and Link State Algorithm require each node maintains k links delay with its neighbors. Generally speaking, IP-layer network satisfies the condition $k \ll n$, but overlay network can be considered as an approximately fully connected topology (see Fig. 1), namely $k \approx n$. Therefore, it is an impractical idea to directly introduce traditional IP-layer shortest path routing algorithms into large-scale overlay network. The study in [13] shows that it is difficult for a fully connected overlay to support more than 60 nodes when using OSPF-like shortest path routing.

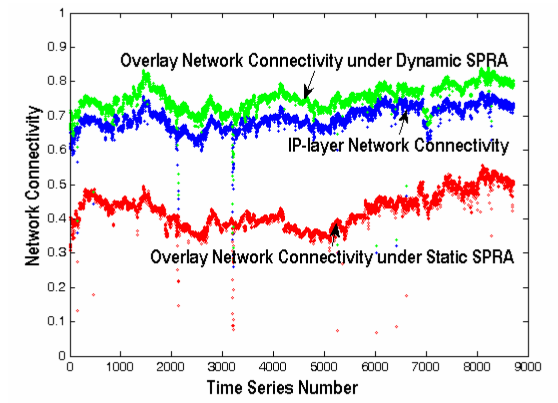

Fig. 6 Network Connectivity Comparison between Static SPRA and Dynamic SPRA

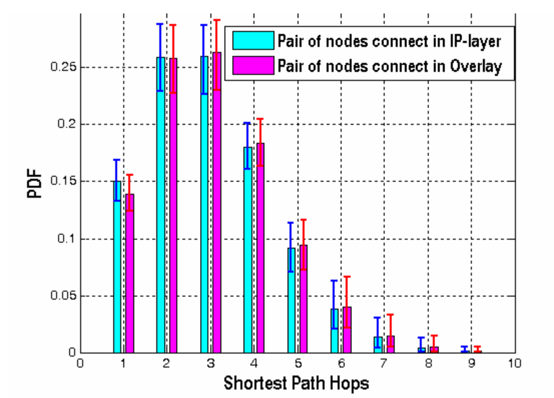

Fig. 7 Error Plot of Shortest Path Hops PDF (with 90% Confidence Interval)

On the other hand, according to Fig. 7, the Error Plot of shortest path hops probability distribution function (PDF) getting from dynamic SPRA, 84% of the shortest paths have less than 4 hops. The average hop of all the shortest paths is 3. This enlightened us to explore 2-Hop overlay routing to lighten the maintenance cost for each node and transmission load for the whole network.

Under 2-Hop overlay routing algorithm (2-Hop ORA), each node just needs to maintain its IP-layer RTT with other nodes and update the RTT change locally, routing table exchanging in Distance Vector Algorithm and update message flooding in Link State Algorithm are avoided. Overlay path searching process is not proactive but reactive, just exchanging locally maintained RTT table between the source node and destination node. The evaluation of 2-Hop ORA will be shown in the following section.

C Multiple Paths Overlay Routing

Compared with traditional IP-layer routing, routing through application layer has high flexibility to support client's diversiform QoS requirements, while is unavoidable to encounter with low reliability rising from the frequent joining and exiting of the relay node. So it is more reasonable to communicate through overlay path and IP-layer path at the same time to guarantee reliability. Besides, E2E delay will become the minimum of the two path RTTs under this multiple paths scheme. Next, we evaluate the performance of 2-Hop overlay routing and multiple paths overlay routing by comparing with IP-layer routing and Dynamic SPRA in three

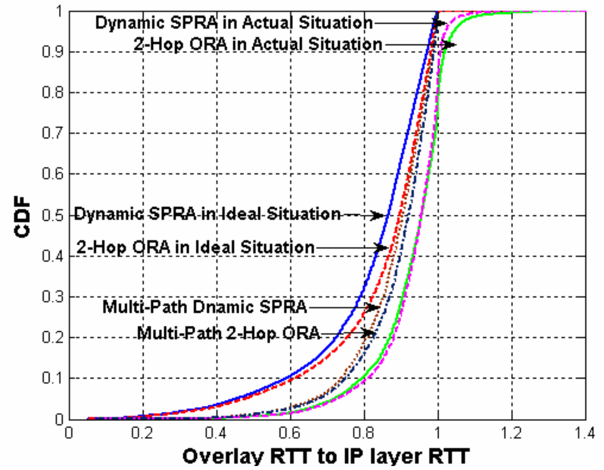

Fig. 8 CDF Curve of Overlay RTT to IP-layer RTT in 3 Different Situations

different situations respectively.

First, ideally, every node knows its current RTT to other nodes accurately. Then the WDG for finding overlay path in each interval is equivalent to the Real-time WDG defined in Section III. However, as already shown in Fig. 2, E2E delay varies rather significantly. It is arduous and costly to get accurate RTT between each pair of nodes. So this comparison just indicates upper bound improvement of the 2-Hop ORA and Dynamic SPRA.

Second, take actual situation into consideration, using the minimum RTT of the latest two intervals to estimate currently RTT [7]. Then the WDG for finding overlay path in the ith interval can be obtained from the $(i-2)$th and the $(i-1)$th Real-time WDG; the IP-layer RTT for comparison is still obtained directly from the ith Real-time WDG the same as the first situation.

Third, two terminal nodes communicate through the found overlay path in the second situation and current IP-layer path at same time (multiple paths).

It can be concluded from Fig. 8 that 2-Hop ORA can reduce majority of E2E delay in any cases. Though in ideal situation, the CDF curve of 2-Hop ORA has some gap to the Dynamic SPRA, these two algorithms can hardly be distinguished in the two actual situations.

As network connectivity achieves its upper bound by Dynamic SPRA, $\delta = \frac{Connect\ Intervals\ by\ ORA}{Connect\ Intervals\ by\ Dynamic\ SPRA}$ can be employed for each pair of nodes to qualify the connectivity performance of 2-Hop ORA. Fig. 9 shows that, for about 10% pair of nodes, δ is under 70% by just IP-layer routing, which means these pair of nodes cannot communicate in over 30% intervals compared with dynamic SPRA. On the other hand, δ can achieve above 90% for over 95% pair of nodes by 2-Hop overlay routing, and multiple paths routing is superior to single 2-Hop routing.

D Discussion

It should be pointed out that by analyzing 96 archives in August 25th, 2005, we found 3-Hop overlay routing can merely save 5ms than 2-Hop overlay routing with the expense of rather heavier calculation load. So it could be concluded

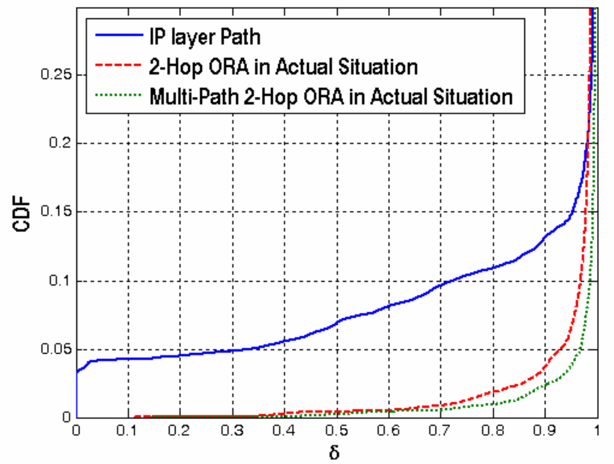

Fig.9 Connectivity Comparison

that only one relay node can efficiently and sufficiently leverage the benefits of overlay routing.

Accurate experiment on 3 PC with Intel Pentium IV 2.40G CPU shows forwarding process added by the relay node takes less than half a millisecond even the relay node's CPU is under rather high load, which proves that the process delay spent in up and down the relay node's stack is negligible compared with the E2E delay mainly varies from dozens to hundreds milliseconds.

V RELATED WORK

One of the early uses of overlay network in the Internet was EON (Experimental OSI-based Network) [14], which proposed an overlay on top the IP network that would allow experimentation with the OSI network layer. The scheme was only experimental and did not lead to a practical deployment for new services or protocols. Then the deployment of multicast protocol took place on an overlay called the MBone [15], but overlay was not generally regarded as a distinct area of research on their own. Later the X-Bone [16] system was built which uses a graphical user interface for automatic configuration of IP-based overlay network.

Detour project [17] at Univ. of Washington was one of the earliest to observe the potential for performing E2E alternate path routing on the Internet. Savage et al. in [18] demonstrated that path selection in the wide-area Internet is sub-optimal from the standpoint of E2E delay, packet loss rate, and TCP throughput. Using dataset of 20 to 40 nodes, they found that for roughly 10% of the host pairs, the best alternative has 50% lower latency.

MIT Resilient Overlay Networks (RON) [19] was a famous overlay network creates a fully connected graph between several nodes, monitors the connectivity between them, and, in case of Internet route failures, redirected packets through alternate overlay nodes. Andersen et al. in [13] found that 51% of the time, improved delay could be obtained via the overlay. They also found a single-hop indirection to be sufficient. But the analogical OSPF overlay routing limits their overlay scale to just dozens of nodes.

Berkeley OverQoS [20] proposed an overlay protocol that uses both retransmissions and forward error correction (FEC) [21] to provide packet loss and throughput guarantees. OverQoS depended on the existence of an external overlay system (the authors suggest RON as an option) to provide path selection and overlay forwarding.

Recently, some study implement overlay network to provide new services. Amir et al [22] used the open source Spines [23]overlay network to segment E2E paths into shorter overlay hops and attempts to recover lost packets using limited hop-by-hop retransmissions. The results showed that Spines can be very effective in masking the effects of packet loss, thus offering high quality VoIP even at loss rates. Overphone [24] re-routesd VoIP calls through overlay network and improvesd the quality of more than a third of the paths when using the G.711 codec. DONet [25] was a Data-driven Overlay Network for live media streaming. The experiments on PlanetLab showed that DONet achieved quite good streaming quality, low level control overhead and transmission delay.

VI CONCLUSION AND FUTURE WORK

This paper quantitates the E2E delay benefits that can be gained as performance improvement from current Internet through overlay routing based on the PlanetLab APP dataset. It is found that the RTT between pairwise PlantetLab nodes could be reduced by 40ms and the network connectivity could be increased by 7% on average under ideal overlay routing algorithm. In-depth investigation indicates that a dynamic, 2-Hop overlay path plus IP-layer path can best leverage the benefits.

The preliminary results of the research presented in this paper also suggest several directions that warrant further research:

First, the 2-Hop ORA presented in this paper still requires each node maintains its RTT with all the other nodes and may suffer a lot with the scale up of overlay network. So it is worthwhile and challenging to design a more scalable and effective overlay routing algorithm to manage and organize the overlay nodes.

Second, consider historical and geographical factors to help choosing better relay nodes. For example, select the node with lower E2E delay and better availability in history with higher possibility; nodes geographically close to each other are inclined to have similar delay characteristics.

Third, the bandwidth, jitter, and packet loss can also be part of parameters to form the weighting of directed edge besides RTT. Thus the original overlay routing algorithm could be extended to support various QoS requirements including E2E delay. It is a prospective direction for future work to design adaptive routing scheme according to individual application requirements.

ACKNOWLEDGMENTS

The work is supported by NEC Laboratories China. The authors would like to thank Yong Xia for his helpful comments, and are grateful to Jeremy Stribling for sharing the data and providing some extremely detailed explanations.

REFERENCE

[1] VoIP-Info, "http://www.voip-info.org/wiki/."

[2] T. Henderson and S. Bhatti, "Networked games - A QoS-sensitive application for QoS-insensitive users?" *Proceedings of the ACM SIGCOMM Workshops*, Karlsruhe, Germany, pp. 141-147, 2003.

[3] W. Jiang and H. Schulzrinne, "Comparison and optimization of packet loss repair methods on VoIP perceived quality under bursty loss," *Proceedings of the International Workshop on Network and Operating System Support for Digital Audio and Video*, Miami, Florida, United States, pp. 73-81, 2002.

[4] C. A. Eldering, M. L. Sylla, and J. A. Eisenach, "Is there a Moore's Law for bandwidth?" *IEEE Communications Magazine*, vol. 37, pp. 117-121, 1999.

[5] All-Pairs-Pings, "http://pdos.csail.mit.edu/~strib/pl_app/."

[6] PlanetLab, "http://www.planet-lab.org."

[7] L. Tang, H. Zhang, J. Li, and Y. Li, "End-to-End Delay Behavior in the Internet," *Proceedings of the 14th IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2006)*, Monterey, California, United States, 2006

[8] S. Warshall, "A Theorem on Boolean Metrices," *Journal of the ACM*, vol. 9, pp. 11~12, 1962.

[9] R. W. Floyd, "Algorithm 97 (SHORTEST PATH)," *Communications of the ACM*, vol. 5, pp. 345, 1962.

[10] C. J. Bovy, H. Mertodimedjo, G. Hooghiemstra, H. Uijterwaal, and P. V. Mieghem., "Analysis of End-to-end Delay Measurements in Internet," *Proceedings of the Passive and Active Measurement Workshop （PAM,2002）*, Fort Collins, Colorado, United States, pp. 26–33, 2002.

[11] S. Gerding and J. Stribling, "Examining the Tradeoffs of Structured Overlays in a Dynamic Non-Transitive Network," Class project: http://pdos.csail.mit.edu/~strib/docs/projects/networking_fall2003.pdf, 2003.

[12] H. Rahul, M. Kasbekar, R. Sitaraman, and A. Berger, "Towards Realizing the Performance and Availability Benefits of a Global Overlay Network," MIT CSAIL Report, MIT-LCS-TR-1009, November, 2005.

[13] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris, "Resilient overlay networks," *Operating Systems Review (ACM)*, Banff, Alta., Canada, vol. 35, pp. 131-145, 2002.

[14] R. Hagens, N. Hall, and M. Rose, "Use of the internet as a subnetwork experimentation with the osi network layer," RFC 1070, February,1989.

[15] H. Eriksson, "MBone: the multicast backbone," *Communications of the ACM*, vol. 37, pp. 54-60, 1994.

[16] J. Touch and S. Hotz, "The X-Bone," *Third Global Internet Mini-Conference at Globecom*, Sydney, Australia, pp. 75–83, , November,1998.

[17] Detour, "http://www.cs.washington.edu/research/networking/detour/."

[18] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson, "The end-to-end effects of Internet path selection," *ACM SIGCOMM Comput. Commun. Rev. (USA)*, Cambridge, Massachusetts, United States, vol. 29, pp. 289-99, 1999.

[19] X. Gu, K. Nahrstedt, R. N. Chang, and Z.-Y. Shae, "An overlay based QoS-aware voice-over-IP conferencing system," *2004 IEEE International Conference on Multimedia and Expo (ICME)*, Taipei, Taiwan, vol. 3, pp. 2111-2114, 2004.

[20] L. Subramanian, I. Stoica, H. Balakrishnan, and R. H. Katz, "OverQoS: an overlay based architecture for enhancing Internet QoS," *First Symposium on Networked Systems Design and Implementation (NSDI '04)*, San Francisco, California, United States, pp. 71-84, 2004.

[21] J. C. Bolot, S. Fosse-Parisis, and D. Towsley, "Adaptive FEC-based error control for Internet telephony," *Proceedings of IEEE INFOCOM '99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies*, New York, New York, United States, vol. 3, pp. 1453-60, 1999.

[22] Y. Amir, C. Danilov, S. Goose, D. Hedqvist, and A. Terzis, "1-800-OVERLAYS: using overlay networks to improve VoIP quality," *Proceedings of the 15th International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV 2005)*, Skamania, Washington, United States, pp. 51 - 56, 2005.

[23] Spines, "http://www.spines.org/."

[24] R. K. Rajendran, S. Ganguly, R. Izmailov, and D. Rubenstein, "Performance Optimization of VoIP using an Overlay Network," *Technical report, NEC Laboratories America*, Princeton, New Jersey, United States, 2005.

[25] X. Zhang, J. Liu, B. Li, and T.-S. P. Yum, "CoolStreaming/DONet: A data-driven overlay network for peer-to-peer live media streaming," *Proceedings of IEEE INFOCOM 2005. The Conference on Computer Communications - 24th Annual Joint Conference of the IEEE Computer and Communications Societies*, Miami, Florida, United States, vol. 3, pp. 2102-2111, 2005.