On the Feasibility of Enhancing Interactivity for Real-time Communications using Overlay Routing

Li Tang[1,2], Zhijie Sun[1], Zhen Chen[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
{tangli03@mails., szj03@mails., zhenchen@, junl@}tsinghua.edu.cn

***Abstract*—As real-time communications are developing rapidly on the Internet, overlay routing has been widely proposed as an easily deployable way to meet their stringent requirements of end-to-end (E2E) performance and interactivity. Contrasting to that much previous work focused on the design of novel protocols, some significant issues impacting overlay routing's feasibility of performance improvement and guidelines of practical implementation still lack a careful study. This paper contributes to the understanding of three of such technical issues. The first issue is how long it takes a relay node to forward a packet, and what are the determinant factors. The second one is how long the overlay path (OP) can persist to sustainedly outperforming the direct IP-path (DP), and how to select more persistent OPs. The last issue is that given the existence of performance correlation between E2E paths, whether it is possible to gain benefits by cooperatively manipulating multiple paths.**

Keywords-overlay routing; real-time communications

I. INTRODUCTION

The oncoming prosperity of real-time communications typified by VoIP and interactive video is posing unprecedented challenges to the Internet. Unlike traditional data transmission applications that rarely concern per packet's delay, the systems of real-time communications is highly sensitive to time-limited-efficacy, for a packet arriving later than expected will be useless but wasting resources. Unfortunately, however, the current Internet cannot guarantee E2E performance on the IP-layer. Enlightened by previous work showing that overlay routing can effectively improve E2E availability by detouring IP-layer's path failure [1, 2], researchers have proposed methods using overlay routing to enhance E2E performance as well. Most such work so far has focused on the design and evaluation of protocols by means of offline simulation or analysis of snapshot trace data, while some significant technical issues are blindly overlooked. For example, it is possible that an OP contingently outperforms the DP while the trace data are collected, but the OP is actually incapable of enhancing E2E performance in practice, due to either acute fluctuation of forwarding time or too short persistence.

The main goal of this paper is to contribute to the understanding of three technical issues that significantly affect the feasibility of using overlay routing to enhance interactivity, a general designation of E2E delay and packet loss based on the particular requirements of real-time communications.

First, we study various factors probably affecting the time it takes a relay node to forward a packet. We find that on various implementation platforms, it takes a relay node with typical personal computer's hardware no more than a few milliseconds to forward a packet. Besides, forwarding packets is a CPU insensitive task, which neither burdens the CPU much, nor highly relies its execution time on the CPU load.

Second, we statistically study the persistence of an OP sustainedly outperforming the DP, by analyzing the round-trip-time (RTT) measurements consecutively collected by all pairs pings (APP) project [3] during a whole day. The results show that among all OPs that ever outperforms the corresponding DP at least once, nearly 80% are inferior to the DP more than half the time, and over 50% are unable to sustainedly outperform the DP for a satisfactorily long time. They suggest that the Internet's dynamic behavior makes a large amount of OPs just incidentally outperform the DP, but it is actually ineffective, if not infeasible, to make use of these OPs to enhance interactivity. To deal with this problem, we propose two simple refining approaches to effectively filter out impersistent OPs.

Third, we study finer-grain characteristics of the Internet paths by doing realistic experiments across a small set of PlanetLab nodes [4]. The results implies that in spite of the existence of correlation between Internet's E2E paths, mesh routing [5] and cooperatively manipulating multiple OPs can still further enhance interactivity than if switching all traffic to a single excellent OP as used by most existing protocols.

The rest of this paper is organized as follows. Section II reviews related work. Section III discusses the time it takes a relay node to forward a packet under diverse scenarios. Section IV reveals the persistence of superior OPs based on APP's RTT measurements. Section V demonstrates the realistic experiments on PlanetLab and their implications to protocol design of overlay routing. Finally, Section VI summarizes the paper and presents our future work.

II. RELATED WORK

Overlay routing at its earliest was proposed to improve path condition for data transmission on the Internet. RON [2] and SOSR [1] proved that overlay routing could indeed practically improve the E2E reliability, throughput, delay, and loss rate for data transmission. But these protocols are not suitable to real-time communications demanding stringent interactivity.

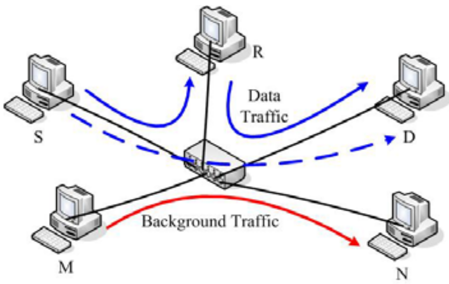

Figure 1. Sketch map of the test bed.

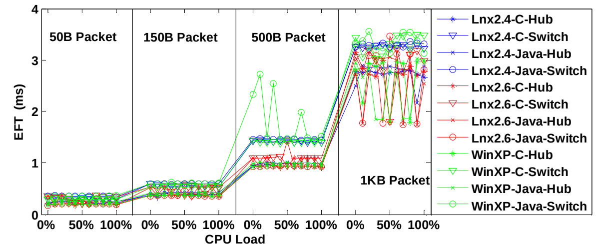

Figure 2. EFT without background traffic under diverse scenarios.

Motivated by the challenge of providing high-quality voice service on the Internet, recent research has devoted to overlay routing schemes specifically for VoIP systems [6, 7, 8]. They mainly focused on designing architectures and protocols for efficiently selecting excellent OPs based on active probing or Internet's topological heuristics, without carefully considering the forwarding time introduced by relay nodes. In this paper, we study the time it takes a relay to forward a packet under diverse scenarios, and provides guidelines of overlay routing design for both VoIP and other real-time communications.

Much work has revealed that the current Internet's routing infrastructure is far from providing the best E2E performance [9, 10, 11], and in considerable part of the cases, there is an alternate OP significantly superior to the DP [12]. Nevertheless, these findings is insufficient to paved the way for the feasibility of using overlay routing to enhance interactivity, as they have neglected a significant issue as discussed in Section 4 of this paper, i.e. whether the utilized OP can be superior persistently enough to pay off the overhead maintaining cost.

In this paper, we limit our discussion to 2-hop OPs, based on previous research showing that most benefits of overlay routing can be obtained by utilizing just a singular relay node [13, 14]. In the rest of this paper, we refer to 2-hop OP as OP for short, except explicit statement is necessary in the context.

This paper is also related to the particular requirements of real-time communications on E2E performance. Specifically, we borrow lessons from ITU standard G.114 [15], which assesses the loss of interactivity because of delay in VoIP systems, declaring that one-way delay up to 400 ms is acceptable for planning purposes, but not for highly interactive conversations, while 150 ms satisfies most applications.

III. FORWARDING TIME OF PACKETS

A. *Test Bed*

Fig. 1 demonstrates the test bed used for measuring the forwarding time, where five computers each with a 100Base-T Ethernet adapter are connected together with a 100Base-T switch or a 10Base-T hub; S , D , and R respectively play the roles of sender, receiver, and relay node, while M and N control background traffic with Iperf [16]. The relay node R is equipped with an AMD Sempron 2800+ processor (1.60GHz) and 256MB memory. In regard to implementation platforms, we examine three operating systems (OS) and two programming languages (PL), as summarized in Table I.

Given an implementation platform, the forwarding time is measured as follows. First, S sends 100 packets to D with an interval of 50 ms. Every packet is duplicated into two copies sent back-to-back through different paths, one directly and the other relayed by R . During this procedure, S and D respectively record the sending gap and arriving gap between the direct and relayed copies of each packet. Then, each packet's forwarding time equals to its arriving gap subtracting its sending gap. Note there is no need to concern the time lag between S 's and D 's clocks, as the forwarding time is only dependent on the difference between the delay of two paths. Finally, the median of the 100 packets' forwarding time is considered as their mathematical expectation (as the mean is easily biased by a few abnormally large flying points), referred to as expected forwarding time (EFT) in the rest of this paper.

TABLE I. IMPLEMENTATION PLATFORMS

Symbol		**Detailed Info**
OS	WinXP	Windows XP Professional; Version 5.1.2600 SP2
	Linux2.4	Redhat 9; Kernel Version 2.4.20-8
	Linux2.6	Fedora Core 5; Kernel Version 2.6.15-1.2054_FC5
PL	C	Visual C++ .NET 2003 (WinXP); GCC 3.22 (Linux2.4); GCC 4.0 (Linux2.6)
	Java	1.5

B. *Results of Experiments*

We first test EFT without any background traffic to study the effect of CPU load and implementation platform. In this test, no packet loss occurs at all. Fig. 2 shows EFTs under different CPU load on diverse implementation platforms. It reveals that EFT appears to be independent from the relay node's CPU load, indicating that relaying packets is a CPU-insensitive task and thus a typical personal computer is capable of playing the role of a relay node. For one thing, forwarding packets would insignificantly burden the CPU; for another, the frequent variation of a personal computer's CPU load would rarely flutter the OP's delay. Although it is noticeable that EFT with the 1KB packets varies much more severely than that with smaller packets, the variation is still limited to no more than 2ms, quite tiny compared with the variation caused by network dynamics. These findings extinguish uncertainties on the feasibility to utilize the common user's personal computers as relay nodes in a large-scale P2P network having a community of millions of users, such as Skype.

Fig. 2 also indicates that EFTs roughly follow a linear relationship with the packet's (payload) size, implying that the forwarding time mainly consists of two parts, the transmission time determined by packet size plus a constant time for packing, unpacking, and header processing. Moreover, as the switch

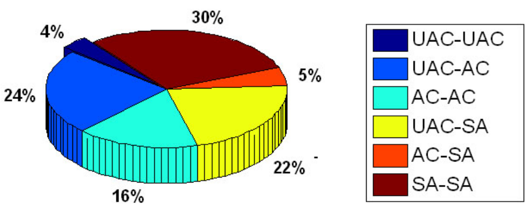

Figure 3. Composition of OP-superior comparisons, where the label indicates the interactivity level of the DP and OP of each component.

needs to control network traffic flows based on the packet's address information, it leads to larger EFT than if using the hub, with the gap ranging from 0.05 ms to 0.5 ms as the increment of the packet's size. However, our further experiments show that when the background traffic between M and N increases to over half the network device's available bandwidth, the hub suffers extremely large EFT (up to hundreds of milliseconds) and unacceptable packet loss rate (over 90%), while the switch leaves its EFT and packet loss nearly the same as if there is no background traffic at all. This is because the hub is a layer-1 network device with all its ports sharing same conflicting domain, while the switch is a layer-2 device that liberates the data flow between S and D from competing with the background traffic for bandwidth. At last, it is noted that the difference between EFTs on different implementation platforms is negligible. Generally speaking, OS of Linux2.6 and PL of C have a little smaller EFTs than their counterparts.

IV. PERSISTENCE OF OVERLAY PATHS

A. *Background*

In this section, we aim to investigate when an OP is determined to be superior to its corresponding DP at present, how long its performance advantage can persist in future. To this end, given a series of sampling measurement data of a specific overlay network along a contiguous period of time, we define and statistically study two metrics for all the OPs that ever outperform their corresponding DPs at least once. One metric is effective percentage (EPG), defined as the proportion of the number of samples in which the OP is superior to the number of all samples. The other is so-called average superior duration (ASD), where a superior duration is defined to be the longest period of time in which all samples exhibit that the OP outperforms the DP. While EPG only concerns with the number of the special samples, ASD is also related to their distribution, or more specifically divergence.

The following investigation in this section is based on APP's RTT data, which are measured with a kernel 'ping' and archived with an interval of 15 minutes. Each archive of the data can be taken as a RTT matrix, in which each element stands for the average RTT between a pair of PlanetLab nodes at that time slot. As PlanetLab nodes have quite special characteristics of joining and leaving, we select the intersection of node sets in all intervals, in order to prevent effect of the churn of overlay network. Besides, we further exclude a few other special nodes to prevent interference of possible node failures during the data collection. Specifically, we base our ensuing analysis on RTT matrices of 96 consecutive intervals during the whole day of November 30, 2005. After the above pretreatment, the final overlay network constantly consists of 188 nodes, forming 35344 node-pairs.

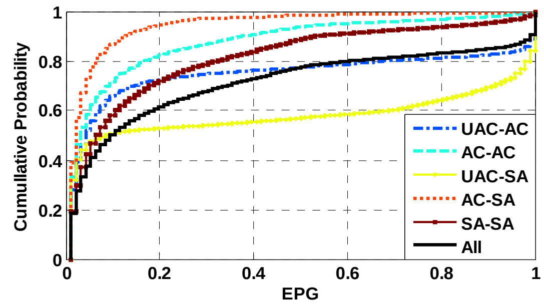

Figure 4. CDFs of the EPGs of SOPs in different categories.

The following terminology is used for clear elaboration: $DP(s,d)$ stands for the DP from source node s to destination node d; $OP(s,r,d)$ stands for the corresponding OP with r as the relay node; $RTT^{\alpha}(s,d)$ stands for the average RTT of $DP(s,d)$ in the αth interval, while the average RTT of $OP(s,r,d)$ in the same interval is defined to be $RTT^{\alpha}(s,r,d) = RTT^{\alpha}(s,r) + RTT^{\alpha}(r,d)$. The OP $OP(s,r,d)$ is considered to be superior to corresponding DP $DP(s,d)$ in the αth interval, if and only if $RTT^{\alpha}(s,r,d) < RTT^{\alpha}(s,d)$. If an OP is ever superior to corresponding DP in at least one interval, the OP will be referred to as a SOP.

Assuming E2E delay is symmetric, and taking RTT as double of one-way delay, we classify the interactivity of a path into three levels by comparing its RTT with the thresholds given by ITU-T G.114, as shown in Table II.

TABLE II. DEFINITION OF RTT-BASED INTERACTIVITY LEVELS

Interactivity level	Range of one-way-delay / RTT (ms)
Satisfied (SA)	$[0,150)$ / $[0,300)$
Acceptable (AC)	$[150,400)$ / $[300,800)$
Unacceptable[a] (UAC)	$[400,\infty)$ / $[800,\infty)$

a. including unavailable

Out of all the ($188^3 \times 96 = 637888512$) comparisons between every OP's and corresponding DP's RTTs in every interval, around 11% are the cases that the OP is superior. This OP-superior portion can be further classified into six different categories, as shown in Fig. 3, where each category represents different worthiness of the interactivity enhancement. For example, 'UAC-UAC' means that although using the OP indeed reduces E2E delay than if using the DP, it is actually useless for enhancing interactivity of real-time communications, as the interactivity level of both the DP and OP is unacceptable. The following study, therefore, excludes the category of 'UAC-UAC' by default. For the other categories, the reduction of delay is practically effective for enhancing interactivity, but it still depends on whether the OP can persist to be superior for a reasonably long time, as to be studied in the rest of this section.

B. *Effective Percentage of SOPs*

Intuitively, the larger EPG a SOP possesses, the more persistent it is, and the more feasible it can be utilized for enhancing interactivity. The plot labeled 'All' in Fig. 4 illustrates the cumulative distribution function (CDF) of all SOP's EPGs, showing that nearly 80% SOPs have their EPGs smaller than 0.5, in other words, only a little more than 20%

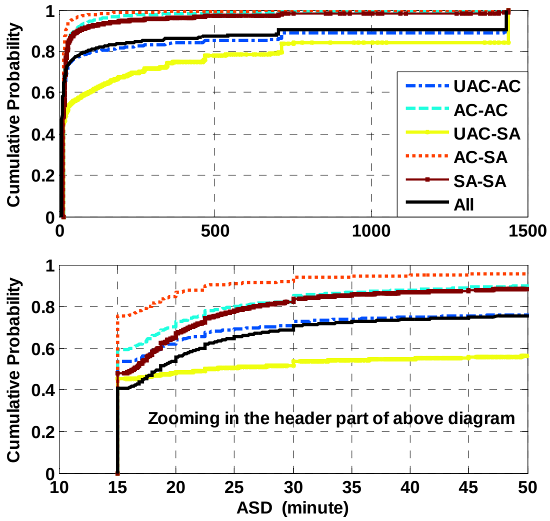

Figure 5. CDFs of the ASDs of SOPs in different categories.

SOPs can be used to replace the DP and statistically enhance interactivity. On the other hand, however, around 60% SOPs have too small EPGs (less than 0.2) to repay their maintenance cost and other overhead (such as the increment of bandwidth requirement), when making use of them for enhancing interactivity of real-time communications.

Fig. 4 also illustrates the EPG CDFs of SOPs in different categories. It is interesting to point out that the category of SOPs with its enhancement of interactivity most meaningful, the one labeled as 'UAC-SA', statistically has the greatest EPGs, which can be considered as a positive evidence for the power of overlay routing techniques. It is worth noting that an SOP is possible to be included in the statistics of more than one categories, because it can be in one category sometime, but in other categories at other times.

C. *Average Superior Duration of SOPs*

Given an interval is 15 minutes, Fig. 5 presents the CDFs of all SOPs' ASDs. As can be seen, nearly 60% SOPs have ASDs larger than 20 minutes, while more than 20% have their ASDs of several hours. Moreover, there are around 10% SOPs having ASDs equal to the whole investigated time (24 hours), meaning they are likely to be able to consistently enhance interactivity over corresponding DPs in all possibility. In respect to different categories, SOPs in the most meaningful category labeled 'UAC-SA' once again statistically have larger ASDs than those in any other category, similar to the situation in terms of EPGs.

We also note an evidently positive correlation between a SOP's EPG and ASD, as is suggested by the similarity between the plots (of the same category) in Fig. 4 and Fig. 5. That is, a SOP having large EPG is likely to have large ASD as well, and vice versa. This indicates that there are indeed some SOPs persistently superior to the DPs, and thus feasible to be used for enhancing interactivity. Note the correlation between a SOP's EPG and ASD is not necessary to exist without verification, as theoretically it is possible for a SOP to simultaneously have large EPG (ASD) and small ASD (EPG). For example, if a SOP and its corresponding DP alternately outperforms each other, the SOP's EPG is 0.5, but its ASD is only 15 minutes.

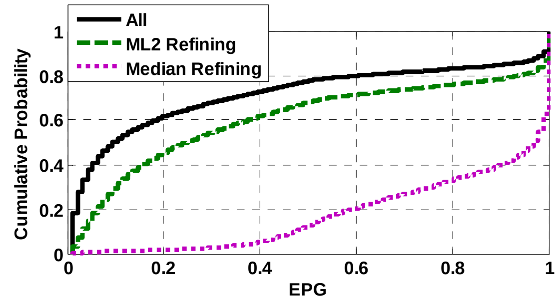

Figure 6. CDFs of the EPGs of all SOPs with/without refinement.

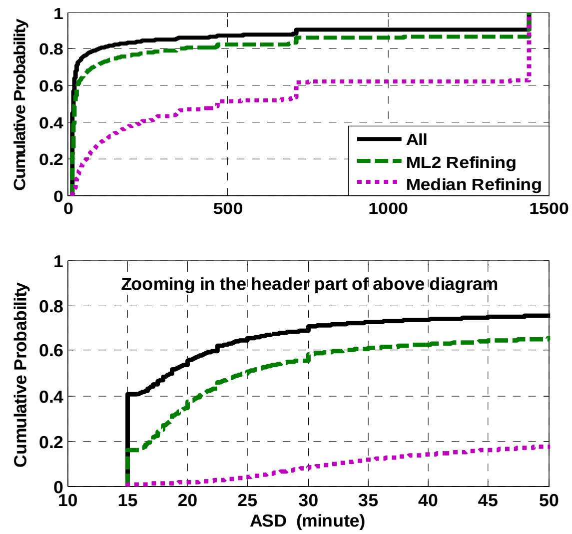

Figure 7. CDFs of the ASDs of all SOPs with/without refinement.

D. *Approaches for Filter out Impersistent SOPs*

As shown above, although around 20% SOPs are persistent enough to feasibly enhance interactivity, such a low proportion makes it rather difficult to effectively select them out. To solve this problem, we propose and examine two simple approaches to filter out impersistent SOPs. One is to determine SOPs based on the so-called min-latest-2 (ML2) RTT matrices, in which the RTT from s to d in the αth interval equals to the minimum of the latest two RTT measurements, $RTT^{\alpha-1}(s,d)$ and $RTT^{\alpha}(s,d)$. This approach comes from observation in our previous research that using the minimum of the most recent two measurements to predict near-term RTTs is preferably accurate [17]. The other approach is to determine SOPs based on the median RTT matrix, of which the RTT from s to d is defined to be the median of the RTTs from s to d in all observed intervals. This approach is based on that RTTs along a reasonably long term generally follow a Gamma-like distribution whose mathematical expectation can be accurately estimated by the median of long-term measurements [18].

Fig. 6 and Fig. 7 respectively present the EPGs and ASDs of all SOPs with/without refined by each approach proposed above. As can be seen, both approaches are effective for filtering out impersistent SOPs. The effect of the second approach based on the median RTT matrix is extraordinarily significant, not only increasing the proportion of SOPs with their EPGs larger than 0.5 from less than 20% to more than

TABLE V. SUMMARY OF THE SELECTED NODE-PAIRS

Symbol	Type	Src	Dst	Average RTT of DP (ms)
US-CN	Cross continents	planetlab-9.cs.princeton.edu	zju1.6planetlab.edu.cn	306.26
CN-BR	Cross continents	tongji2.6planetlab.edu.cn	Planetlab2.lsd.ufcg.edu.br	486.15
CN-CN	Domestic	xmu1.6planetlab.edu.cn	Csu2.6planetlab.edu.cn	66.00
US-US	Domestic	planetlab8.millennium.berkely.edu	planetlab3.csail.mit.edu	93.04

80%, but also making more than 80% of the refined SOPs have ASDs over 50 minutes. It means after such a refinement, even randomly selecting from the remaining SOP candidates, it is still highly probable to pick a SOP that is both statistically superior to the DP and persistent enough for practical usage.

To examine the impact of these two refining approaches on the decrement of SOPs, we count the number of SOPs from three different viewpoints, namely the number of all SOPs in all intervals (#All-SOP), the number of distinct SOPs in all intervals (#Distinct-SOP), and the average number of SOPs between a pair of nodes in an interval (#Avg-SOP). As shown in the first two columns of Table III, the two approaches respectively reduce #All-SOP by 6.9% and 27.0%, #Distinct-SOP by 29.4% and 74.8%, proving that the filtered out SOPs generally have small EPGs. According to the last column, both approaches cause only slight decrement of #Avg-SOP, and can thus still draw out SOP's potential for enhancing interactivity.

TABLE III. #SOPS WITH/WITHOUT REFINEMENT

Improving Approach	#All-SOP	#Distinct-SOP	#Avg-SOP
All without refinement	66686623	2421289	26.05
ML2 RTT matrices based	62080933	1709377	24.98
Median RTT matrix based	48673948	611042	22.06

V. PERFORMANCE CORRELATION BETWEEN OPS AND DPS

A. *Experiment Setup on PlanetLab*

A common concern with multi-path routing is to what extent the paths are correlated to each other. In extremis, multi-path routing would be useless, if all paths share the same bottleneck link. To investigate the performance correlation between an OP and corresponding DP, we conduct experiments emulating the traffic of real-time communications among a small subset of PlanetLab nodes.

The experiments are carried out session by session. Each session consists of 10000 size-fixed UDP packets that are contiguously sent with a fixed interval from a sender node to a receiver node. Similar to the way used in Section III, every packet is duplicated and sent back-to-back through two paths, one is the DP and the other is a specific OP unaltered during the whole session. Every packet contains necessary information for the receiver to identify which session the packet belongs to, what the packet's sequence number is, when the packet is sent and received, and which path the packet passes by. With such information, it is sufficient to figure out a session's loss rate and one-way delay of every successfully received packet.

As the sender's and receiver's clocks are asynchronous, the delay needs to be corrected. To this end, we start another process of 'ping' whenever a session begins to measure RTTs from the sender to receiver. Assuming the average delay of all packets in a session is $AvDLY$ and the average RTT during the session is $AvRTT$, we estimate the time lag between the sender's and receiver's clocks by $T = AvDLY - 0.5AvRTT$ (the routes between our selected nodes are generally symmetric according to the 'traceroute' results). At last, each packet's delay is corrected by subtracting T from it.

Table IV gives the parameters used in the experiment, where proper combinations of the packet size and sending gap roughly emulate traffic patterns of different applications. For example, 150-byte packets with sending gap of 20 ms emulates some PCM-based VoIP traffic, while 500-byte and 5ms may be proximate to the traffic of some high-quality interactive video.

TABLE IV. PARAMETER VALUES USED IN THE EXPERIMENTS

Parameter	Value
Packet[b] size (Byte)	50, 150, 500, 100
Sending gap (ms)	5, 20, 50

b. payload of a UDP packet

To avoid disturbing Internet's normal traffic, we limit our experiments to four node-pairs as summarized in Table V. For each node-pair, three different OPs are surveyed under every combination of the parameters in Table IV, thus making up 144 sessions in total. Specifically, the following study is based on the data collected on December 15, 2006.

B. *Path Correlation and Mesh Routing*

In this section, we study performance correlation between paths through the correlation coefficient of two widely used heuristics for bottleneck links, namely the delay variation and packet loss. The correlation coefficient is defined as

$$\frac{E[XY - E(X)E(Y)]}{\sqrt{E[X^2 - E(X)^2]E[Y^2 - E(Y)^2]}}$$

where $E(f)$ stands for f 's mathematical expectation, and $\{X, Y\}$ stands for the status of every pair of packets sent back-to-back respectively through DP and OP. Specifically, in terms of delay variation, X stands for the delay variation of two neighboring DP-packets, while in terms of packet loss, X has a value of 1 if the DP-packet is lost, or 0 otherwise. Y has the same meaning as X in each case, except for the OP-packets.

Fig. 8 shows CDFs of the correlation coefficient of OP's and DP's delay variation and packet loss. Around 80% (40%) sessions have the OP's delay variation (packet loss) positively correlated to the DP's counterpart. It is likely that the OPs in these sessions share some common bottleneck links as the DP, because the E2E delay variation and packet loss are most usually caused by the appearance or alternation of bottleneck links.

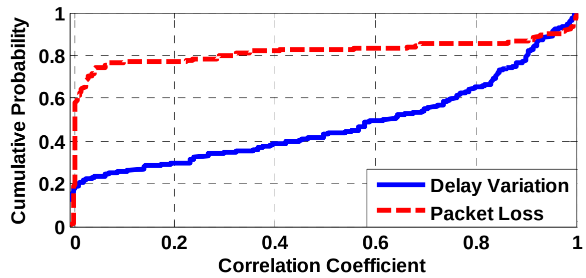

Figure 8. CDFs of the correlation coefficient of OPs and DPs.

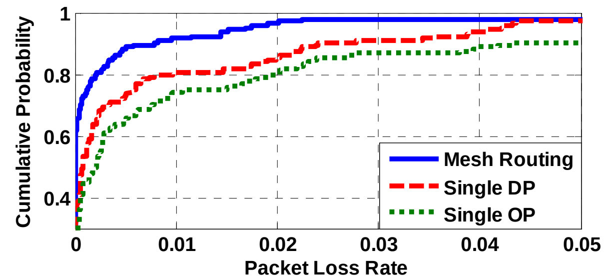

Figure 9. CDFs of the loss rates using mesh routing or single DP/OP.

On the other hand, as quite few OP's delay variation and packet loss are completely predictable by the DP's counterparts (with correlation coefficient being 1.0), it indicates that routing with redundancy is still able to enhance interactivity of real-time communications to some extent, similar to the conclusions of [5] that studies the benefits of multi-path routing for data transmission applications. Fig. 9 illustrates CDFs of all session's loss rates respectively with mesh routing, single DP, and single OP, showing visible improvement of mesh routing on reducing loss rate. In respect to E2E delay, as the small set of OPs are not selected consciously to have shorter delay than corresponding DPs, using mesh routing only reveals trivial advantages over using single DP.

C. Discussion

Another interesting observation based on our experiment results is that a session's E2E loss rate is highly related to its bandwidth requirement and traffic pattern; the E2E delay, however, is nearly independent from these factors. It implies that cooperatively manipulating multiple OPs has potential to further enhance interactivity than using single path, as the former way enables flexibility of adjusting the traffic pattern on each path. We leave detailed investigation to our future work.

VI. CONCLUSION

This paper investigates three significant technical issues on the feasibility of using overlay routing to enhance interactive of real-time communications. The results indicate: A) Relaying packets is a CPU insensitive task: a typical personal computer can relay a packet in no more than a few milliseconds, and the variation of CPU load rarely affects forwarding time. B) In a large-scale overlay network, there are indeed OPs that persistently outperform corresponding DPs, and can be practically utilized; some simple approaches can effectively filter out impersistent OP candidates. C) In spite of the performance correlation between Internet's E2E paths, cooperatively manipulating multiple OPs still provides more flexibility and has potential to further improve than switching all traffic to a single path as favored by most existing research.

While this paper offers some useful design guidelines, our work remains to be done in designing scalable overlay routing protocols and evaluating their effectiveness in practice.

ACKNOWLEDGMENT

This work is sponsored by NEC Laboratories China. The authors are grateful to Hui Zhang and Zhiming Zhang at security Lab RIIT Tsinghua University for their comments.

REFERENCES

[1] K. P. Gummadi, H. V. Madhyastha, S. D. Gribble, H. M. Levy, and D. Wetherall, "Improving the Reliability of Internet Paths with One-hop Source Routing," in Proceedings of 6th Symposium on Operating System Design and Implementation (OSDI), 2004.

[2] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris, "Resilient Overlay Networks," in Proceedings of 18th ACM SOSP Conference, 2001.

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

[4] PlanetLab: http://www.planet-lab.org.

[5] D. G. Andersen, A. C. Snoeren, and H. Balakrishnan, "Best-Path vs. Multi-Path Overlay Routing," in Proceedings of Internet Measurement Conference, 2003.

[6] A. Yair, D. Claudiu, G. Stuart, H. David, and T. Andreas, "1-800-OVERLAYS: Using Overlay Networks to Improve VoIP Quality," in Proceedings of the International Workshop on Network and Operating Systems Support for Digital Audio and Video, 2005.

[7] S. Ren, L. Guo, and X. Zhang, "ASAP: an AS-Aware Peer-relay Protocol for High Quality VoIP " in Proceedings of the 26th International Conference on Distributed Computing Systems (ICDCS), 2006.

[8] S. Tao, K. Xu, A. Estepa, T. Fei, L. Gao, R. Gu´erin, J. Kurose, D. Towsley, and Z.-L. Zhang, "Improving VoIP Quality through Path Switching," in Proceedings of IEEE Infocom Conference, 2005.

[9] V. Paxson, "End-to-End Routing Behavior in the Internet," IEEE/ACM Transactions on Networking, vol. 5, no. 5, pp. 601~615, 1997.

[10] S. Neil, M. Ratul, and A. Thomas, "Quantifying the Causes of Path Inflation," in Proceedings of ACM SIGCOMM Conference, 2003.

[11] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed Internet Routing Convergence," in Proceedings of ACM SIGCOMM Conference, 2000.

[12] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson, "The End-to-End Effects of Internet Path Selection," ACM SIGCOMM Computer Communication Review vol. 29, no. 4, pp. 289-299, 1999.

[13] H. Rahul, M. Kasbekar, R. Sitaraman, and A. Berger, "Towards Realizing the Performance and Availability Benefits of a Global Overlay Network," in Proceedings of 7th Passive and Active Measurement Conference (PAM), 2006.

[14] H. Zhang, L. Tang, and J. Li, "Impact of Overlay Routing on End-to-End Delay," in Proceedings of 15th International Conference on Computer Communications and Networks (ICCCN), 2006.

[15] "One Way Transmission Time," ITU-T Recommendation G.114, 2000.

[16] Iperf: http://dast.nlanr.net/Projects/Iperf/.

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

[18] C. J. Bovy, H. T. Mertodimedjo, G. Hooghiemstra, H. Uiterwaal, and P. V. Mieghem, "Analysis of End-to-End Delay Measurements in the Internet," in Proceedings of the Passive and Active Measurement Workshop (PAM), 2002.