Exploring the Performance Benefits of End-to-End Path Switching*

[Extended Abstract]

Shu Tao, Kuai Xu, Ying Xu, Teng Fei
Lixin Gao, Roch Guerin, Jim Kurose, Don Towsley, Zhi-Li Zhang

Categories and Subject Descriptors: C.2[Computer Systems Organization]:Computer Communication Networks

General Terms: Measurement, Performance

Keywords: Multi-homing, Overlay, Path switching

1. INTRODUCTION

As more applications become distributed in nature, it is increasingly important to ensure the best possible performance when transferring data between different end-points. Furthermore, the increased availability of multiple providers and the development of technologies that let end-users control where and how their traffic is to be sent, make it possible to take advantage of "path diversity" to improve the performance and availability of data transfers for such applications. The ability of a source end system to dynamically switch among multiple paths to a destination is often referred to as "path switching." In this paper, we explore if, when, and how the ability to perform path switching can yield meaningful end-to-end performance improvements.

2. WIDE-AREA MEASUREMENT RESULTS

In order to explore the benefits of path switching, we study the path diversity offered by two common approaches, namely multi-homing and overlay routing, which have been used by both industry and academia [1] [2]. We set up a wide-area testbed, on which (1) each node has the ability to select different providers to reach other nodes, and (2) the nodes can form an overlay network by using another node as an intermediary to reach other nodes. Our testbed involves three multi-homed campus networks in the US, i.e., University of Massachusetts (UMass), University of Pennsylvania (UPenn), and University of Minnesota (UMN). We assessed the level of path diversity by conducting extensive measurements on our testbed across a period of several months to monitor end-to-end *delay* and *loss* characteristics of all possible path combinations.

Based on the measurement, our first major observation is that *for end-to-end delay, there exists a dominant path that almost always outperforms the other paths between a source and a destination node pair.* This is likely because propagation delay is the primary factor that affects the end-to-end delay. The relative delay rankings of parallel paths only change when major events such as AS-level routing changes happen and the durations are usually long. Therefore there is not much delay benefit in performing dynamic path switching, especially at small time scales (e.g., minutes). Long-lived delay performance changes can be easily detected without resorting to any sophisticated prediction mechanism.

On the other hand, there *does not* exist a dominant best performing path with respect to the *end-to-end loss* performance. Most paths experience losses during certain periods of time, and those lossy periods appear sporadic and somewhat uncorrelated among the paths between a source and destination node pair. Such observations suggest that path switching has the potential to improve loss performance. To quantify the performance gains, we consider the *theoretically best attainable* loss performance using path switching by selecting the provider or overlay path with the best average loss rate over each 1-minute interval, assuming that we have *a priori* knowledge of the average loss rate over all the intervals. Our analysis on the collected traces shows that for all source-destination pairs on our testbed, the overall loss rate can be reduced by 64.3%, if path switching is performed at 1 minute time scale. The details of the measurement results and our analysis can be found in [3].

3. THE CASE FOR PATH SWITCHING

Although the theoretical improvement in performance of using path switching is remarkable for end-to-end loss performance, fully realizing this potential calls for an effective method to predict the performance on each available path and make correct path switching decisions.

3.1 Predicting Path States

A random process is predictable only if it exhibits some form of temporal dependency. Our analysis of the traces we gathered shows the presence of temporal correlation and indicates the potential of using past path loss behavior to predict future behavior. For this purpose, we model the path performance using a simple two-state model as follows. Based on the loss rate during an observation time window w, we label each path as being either in *good* or *bad* states: it is in the *good* state (represented by 1) if the average loss

---

[*] Shu Tao, Ying Xu and Roch Guerin are with University of Pennsylvania; Teng Fei, Lixin Gao, Jim Kurose and Don Towsley are with University of Massachusetts, Amherst; Kuai Xu and Zhi-Li Zhang are with University of Minnesota. This work was supported in part by the National Science Foundation under the grants ANI-9906855, ITR-0085930, ITR-0085848, ANI-0073819 and ITR-0085824.

SIGMETRICS/Performance'04, June 12–16, 2004, New York, NY, USA.
ACM 1-58113-664-1/04/0006.

rate in the interval is less than (a pre-specified) threshold θ; otherwise it is in the *bad* (0) state. By examining all the traces, we find that w equal to 1 minute is a good choice due to the strong correlation observed at that lag; $\theta = 3\%$ is a decent choice to define the 2-state mode for prediction as it is a good value for most of our traces.

Given this two-state path performance model, we first study a simple predictor, which always predicts the state in the next interval using the observed path state in the current interval, i.e.,

$$X_i = x_{i-1} \tag{1}$$

Alternatively, we also study Markov predictors with various amount of memory, namely, Markov predictors of different orders. A Markov predictor with an order k is defined as follows: given that the states of a path in the last k intervals, $X_{i-1}, X_{i-2}, ...X_{i-k}$, its state in the next interval i, X_i, is predicted as

$$X_i = \arg\max_{x_i} Pr\{X_i = x_i | X_{i-1} = x_{i-1}, ..., X_{i-k} = x_{i-k}\}$$

Clearly, the order of the Markov model, k, determines how much memory (past history) is needed for predicting the path state in the next interval.

The performance of a predictor is typically measured by two metrics: *precision* rate p – the fraction of predicted states that match the observed states, and *recall* rate r – the fraction of observed states that are correctly predicted. Using these metrics, Fig.1 shows the performance of the simple predictor and the Markov predictors with various orders using one of the datasets. We see that the simple and first-order predictors perform as well as the higher-order predictors. The observation also holds for other datasets.

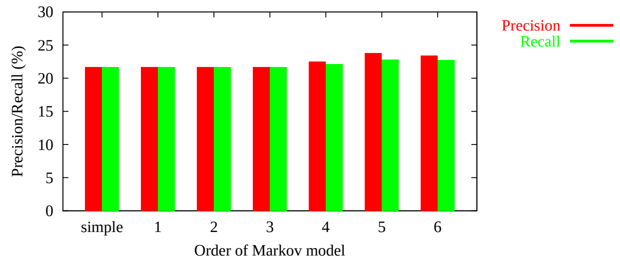

Figure 1: The precision-recall percentage of the simple predictor and Markov predictor with different orders.

As an aside, we also consider "finer-grained" models for the path performance using more than two ("good" and "bad") states. We find that although using a finer definition of path state has the advantage of better distinguishing the quality difference between two paths, the probability of correctly predicting lossy states decreases as the number of states increases.

3.2 Prediction-Based Path Switching

Based on the above analysis, we adopt the following simple path switching strategy: Using the simple predictor (1), we predict the state of each path in the next time interval (1 minute) based on the loss rate measured in the current time interval, and we choose the candidate path with the best predicted state. We continue using the current path unless a better path exists.

To evaluate the performance of this simple path switching mechanism, we use the *theoretically best attainable* loss rate, described in Section 2, as our reference. In Fig. 2, we compare all possible pairs of paths in our testbed, including those available through multi-homing and the use of the overlay network. Each point in the figure corresponds to a different pair of paths. Namely, for a pair of candidate paths with loss rates s_1 and s_2, we first compute the resulting loss rate of prediction-based path switching s', and the relative improvement is then computed as $\frac{\min(s_1, s_2) - s'}{\min(s_1, s_2)}$. The corresponding reference (optimal) value is computed similarly. Fig. 2 shows that for most trace pairs the improvement in loss performance is substantial. Moreover, in some cases the loss improvement achieved by the on-line path switching mechanism is fairly close to the best attainable value.

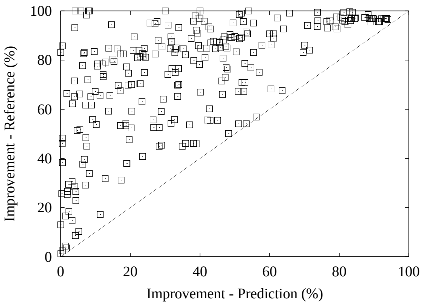

Figure 2: The relative loss rate improvement: path switching versus optimal reference.

4. CONCLUSIONS

In this paper we explored the questions of if, when and how path switching is beneficial in improving end-to-end performance. For this purpose we set up a wide-area multihoming testbed with rich path diversity afforded by both multiple service providers as well as overlay paths. Using the delay and loss measurement traces collected on this testbed, we investigated the potential benefits of dynamic end-to-end path switching in terms of both delay and loss performance. Our conclusions are two-fold: 1) While there are significant differences in end-to-end delay performance among various paths, such differences are in general fairly static (i.e., long lasting) and thus can be utilized in a straightforward manner. 2) Whereas, loss rates of different paths vary dynamically, and there is *no* single "dominant" path that always outperforms others. Hence *dynamic* path switching among different paths can yield potential performance benefits. Furthermore, we demonstrated that such performance benefits can be effectively exploited by using a simple dynamic path switching strategy.

5. REFERENCES

[1] RouteScience. http://www.routescience.com/.
[2] D. G. Andersen et al. Resilient Overlay Networks. In Proc. of SOSP, October, 2001
[3] S. Tao et al. Exploring the performance benefits of end-to-end path switching. Technical report, available at: http://www.cs.umn.edu/research/networking/itr-qos/papers/tech-report.pdf.