Building and maintaining overlay networks for bandwidth-demanding applications

Min Sik Kim

Advisor: Simon S. Lam

Department of Computer Sciences, The University of Texas at Austin

Introduction

As the Internet grows in scale, the demands of applications also grow in terms of required resources and types of services. Overlay networks have emerged to accommodate such applications. The performance of an overlay network is, however, highly dependent on its topology; keeping logical links in an overlay network from interfering with each other is crucial in achieving high throughput.

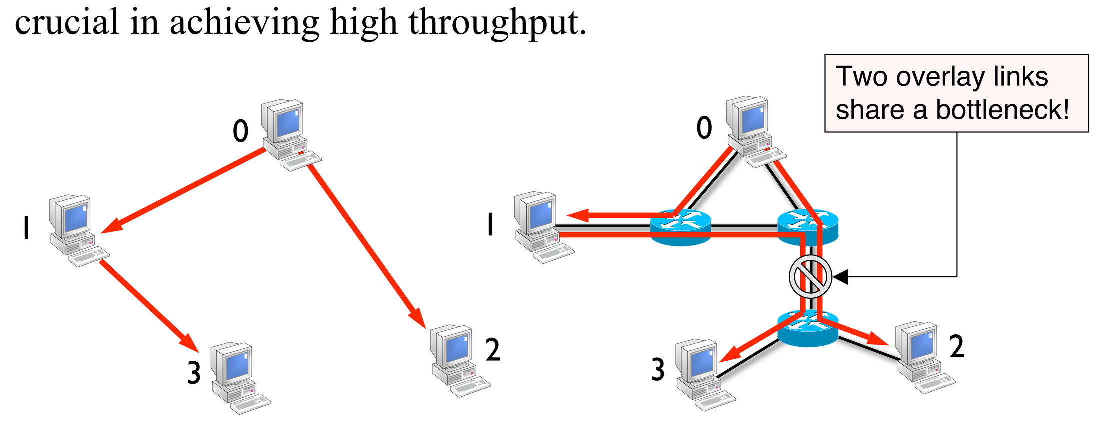

Overlay networks create a bottleneck because multiple overlay links share the same physical link.

Shared bottleneck detection

A key observation is that the delay of a congested link (bottleneck) has large fluctuations due to queueing delay changes, while the delay of a link with light load is relatively stable. Therefore, measured delays of two paths show strong correlation if the paths share one or more congested links, and little correlation if they do not share any congested links [RKT02]. The cross-correlation coefficient (*XCOR*) is used to measure such correlation.

$$XCOR_{XY} = \frac{\sum_{i=1}^{n}(X_i - \overline{X})(Y_i - \overline{Y})}{\sqrt{\sum_{i=1}^{n}(X_i - \overline{X})^2 \cdot \sum_{i=1}^{n}(Y_i - \overline{Y})^2}}$$

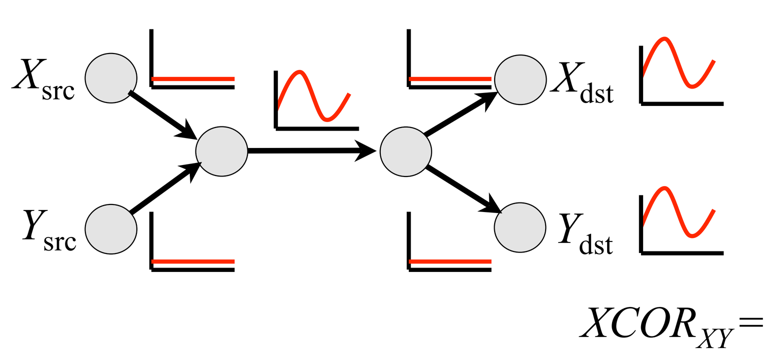

$XCOR_{XY} = 1$

XCOR is 1 If paths X and Y share congested links.

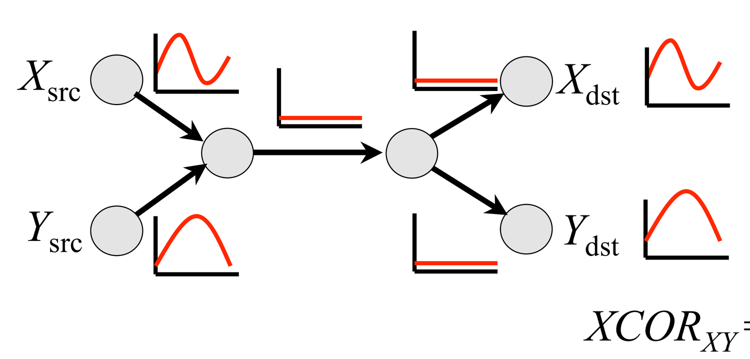

$XCOR_{XY} = 0$

XCOR is 0 if paths X and Y do not congested links.

Limitations

Applicability of the simple *XCOR* technique is limited because it makes two assumptions that generally do not hold for the Internet.

- Queueing delay variation on non-congested links is close to zero.
- The two delay sequences are synchronized.

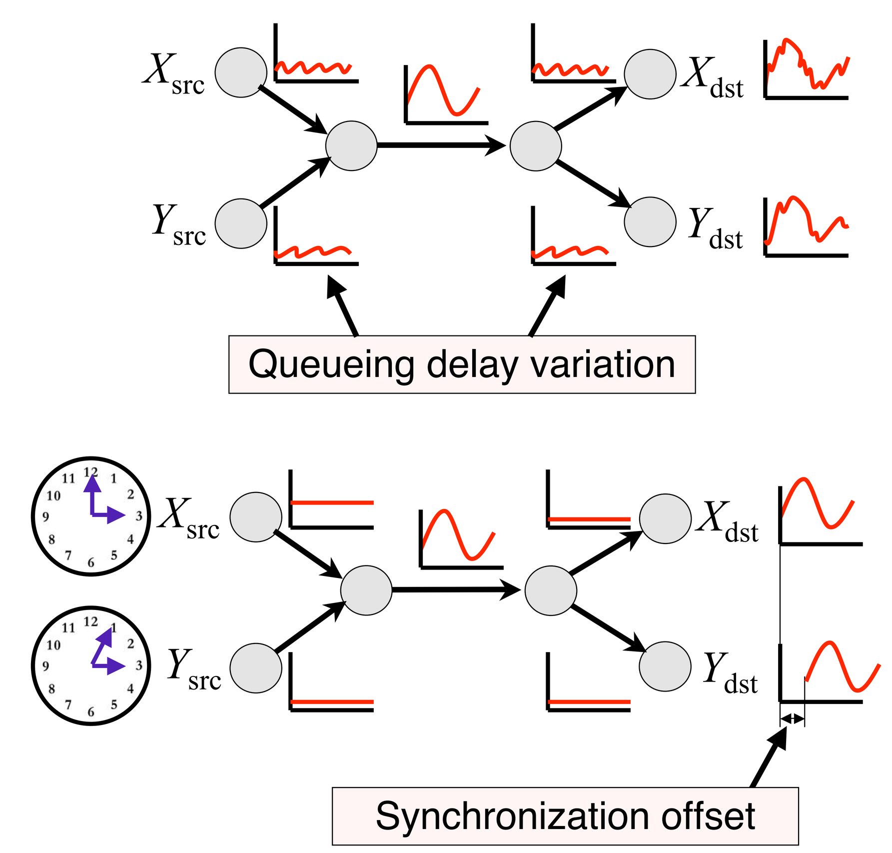

Wavelet-based approach [KKS+04]

The wavelet transform is a signal processing technique that represents a transient or non-stationary signal in terms of time and scale distribution. Due to its light computational complexity, the wavelet transform is an excellent tool for on-line data compression, analysis, and denoising. In particular, wavelet denoising is capable of separating random queueing delay variation from congestion-induced delay.

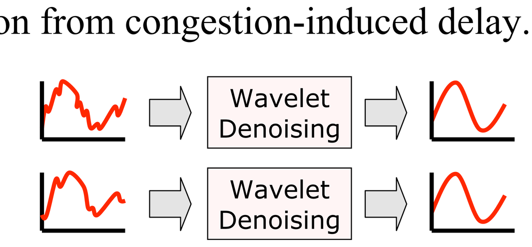

The time and scale resolution of the representation depends on the selection of a wavelet basis. By carefully choosing the basis function, the adverse effect of synchronization offset can be minimized.

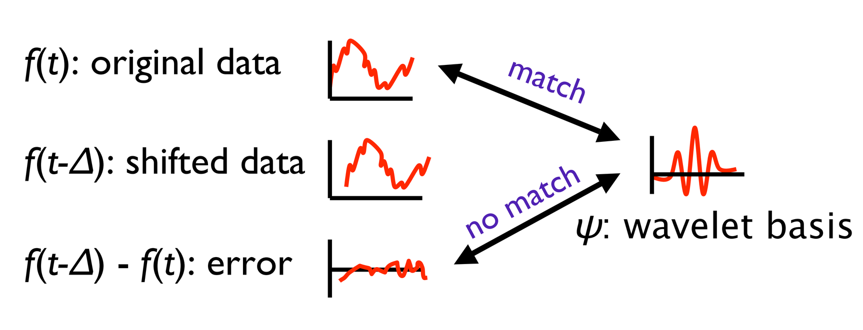

Evaluation

The performance of the wavelet-based technique (DCW, Delay Correlation with Wavelet Denoising) is compared against two representative techniques.

- MP (Markovian Probing) [RKT02]
- BP (Bayesian Probing) [HBB00]

Accuracy of each technique is represented by *Positive Ratio.*

$$\text{Positive Ratio} = \frac{\text{\# of answers indicating shared congestion}}{\text{\# of experiments}}$$

1: shared congestion, 0: no shared congestion

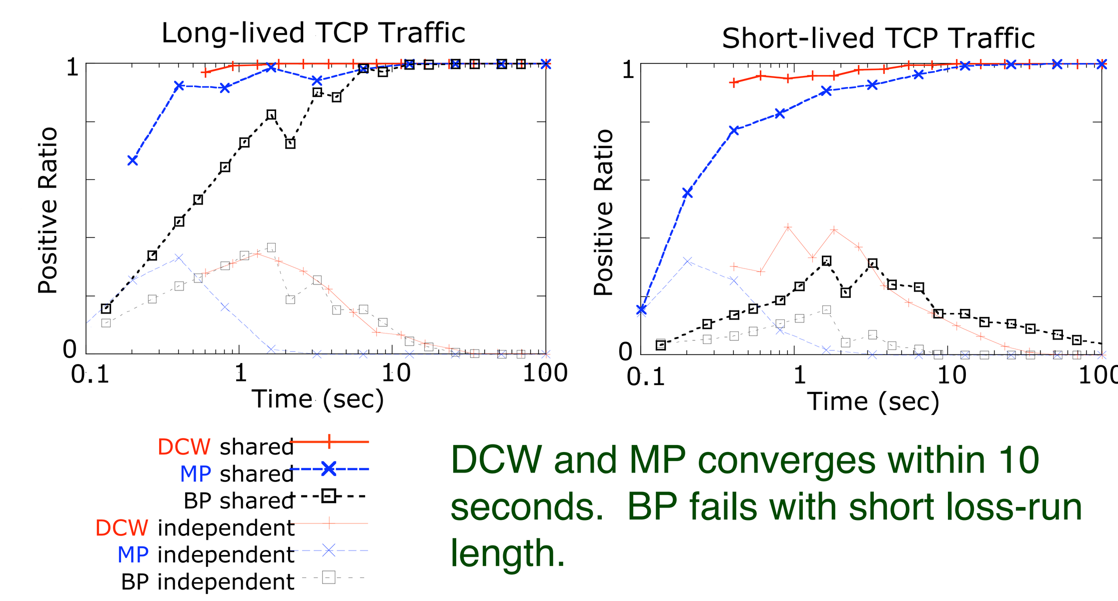

DCW and MP converges within 10 seconds. BP fails with short loss-run length.

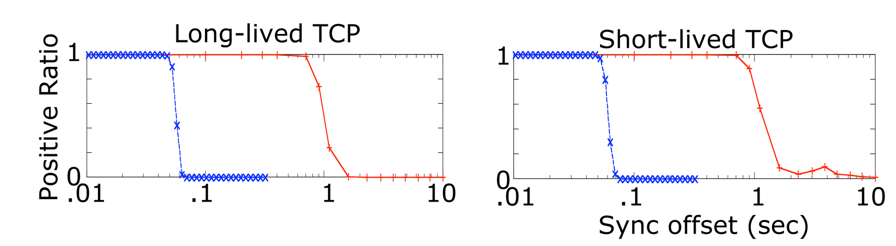

Synchronization offset tolerance

DCW: 1-2 sec, MP: 30-70ms, BP: < 10ms

Scalability

The shared congestion detection technique can be used to find "better" paths in many applications, such as overlay multicast [KLL05], overlay QoS routing, file download from multiple servers, and exploiting path diversity. However, it requires $O(N^2)$ to detect all shared congestion among N paths. Multi-dimensional indexing can reduce the complexity.

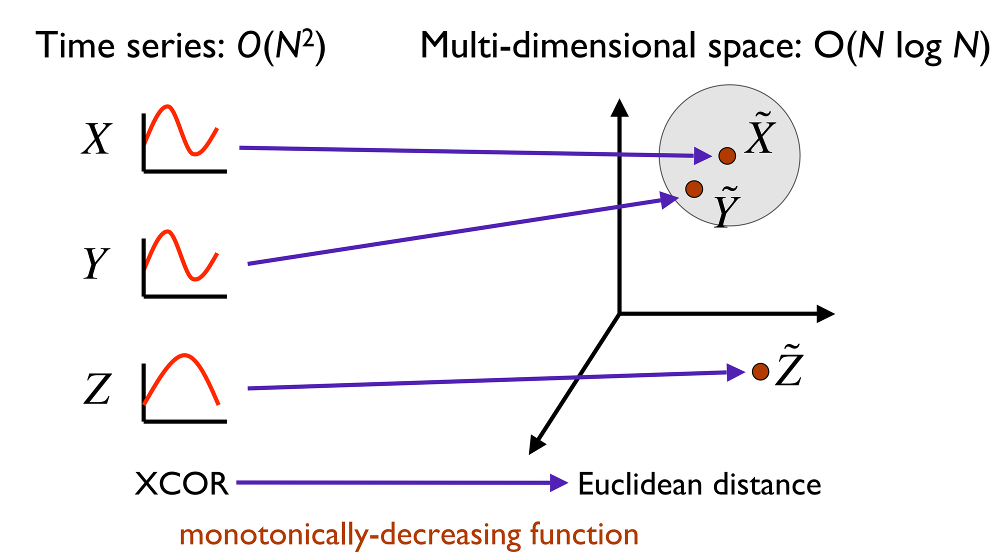

Dimensionality Reduction

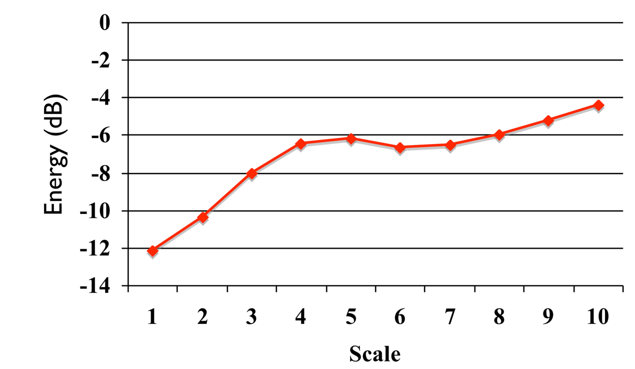

Dimensionality can be reduced by excluding low-energy scales in multi-dimensional mapping.

Conclusions

- Shared congestion detection can identify bottlenecks in overlay networks.
- The wavelet-based approach provides a robust and accurate technique to detect shared congestion.
- Multi-dimensional indexing improves scalability of shared congestion detection.
- This work will serve as a foundation on which future applications can achieve higher throughput by building more efficient overlay networks.

Literature cited

[RKT02] D. Rubenstein, J. Kurose, and D. Towsley. Detecting shared congestion of flows via end-to-end measurement. *IEEE/ACM Transactions on Networking* 10(3):381–395. June 2002.

[KKS+04] M. S. Kim, T. Kim, Y. Shin, S. S. Lam, and E. J. Powers. A wavelet-based approach to detect shared congestion. In *Proceedings of ACM SIGCOMM 2004*. August 2004.

[HBB00] K. Harfoush, A. Bestavros, and J. Byers. Robust identifications of shared losses using end-to-end unicast probe. In *Proceedings of the 8th IEEE ICNP*. November 2000.

[KLL05] M. S. Kim, Y. Li, and S. S. Lam. Eliminating bottlenecks in overlay multicast. In *Proceedings of IFIP Networking 2005*. May 2005.

Acknowledgments

We thank YongJune Shin and Edward J. Powers for advice on wavelet analysis, Taekhyun Kim and Yi Li for experiment assistance, Eddie Kohler for shepherding our early work on shared congestion detection, and anonymous reviews for their constructive comments. Funding for this project was provided by National Science Foundation grants ANI-0319168 and CNS-0434515.

For further information

Please contact *minskim@cs.utexas.edu*. More information on this work can be obtained at *http://www.cs.utexas.edu/~minskim/*.