# Application of Wavelet Denoising to the Detection of Shared Congestion in Overlay Multimedia Networks


Taekhyun Kim*, YongJune Shin*, Edward J. Powers*, Min Sik Kim†, and Simon S. Lam†

*Department of Electrical and Computer Engineering, The University of Texas at Austin
Tel: 1–512–471–1662, Fax: 1–512–471–1856, Email: {thkim,june,ejpowers}@ece.utexas.edu

†Department of Computer Sciences, The University of Texas at Austin
Tel: 1–512–232–7489, Fax: 1–512–471–8885, Email: {minskim,lam}@cs.utexas.edu



***Abstract*— The overlay network approach is an emerging technique to satisfy the strict requirements for various real-time multimedia services. However, overlay networks suffer from a shared congestion problem since each unicast flow may interfere with each other in the common underlying links. Most previous techniques to detect shared congestion have limitations when applied as a general solution, since they assume perfect synchronization between probing packets. However, our recent work shows that a technique based on wavelet denoising can overcome the limitations by mitigating the interfering effects such as synchronization offset and the random fluctuations of queueing delay; the proposed technique provides a more robust and accurate detection in the presence of a large amount of synchronization offset. In this paper, wavelet denoising is tailored to the characteristics of queueing delay on packet networks. The wavelet denoising based technique is verified through extensive simulations. The efficacy of the proposed approach is demonstrated by the detection accuracy and convergence speed.**


## I. Introduction

In recent years, more people are enjoying multimedia services such as audio and video streaming services, remote conference services, etc. This multimedia traffic is growing fast and expected to exceed the traditional Internet service traffic including e-mail and WWW. The transition of Internet traffic from the traditional Internet services to real-time, interactive multimedia services is putting more strict requirements on the underlying networks: broader bandwidth, lower packet loss rate, and delay variance. However, the layered structure of the Internet hinders the introduction of additional features for those requirements, since the improvement in intermediate or core level routers is limited and often impossible. Hence, the approach that application layer hosts process their own routing or forwarding task over the traditional network layer, so called overlay networks, emerges as an alternative way to meet the additional requirements for multimedia services.

In spite of a recent proliferation of overlay networks, many applications of overlay networks, including application layer multicast and file download from multiple hosts, have a congestion problem in the underlying networks. Since there are usually a large number of unicast flows among distributed end hosts, they may interfere with each other by sharing common underlying links. The problem for application layer overlay systems is even worse than the case of network layer systems, because there is no information on the underlying network status. There are already many traditional congestion control approaches. Since the existing congestion control methods are performed on a per-flow base, the knowledge about which unicast flows are sharing a congested underlying link can be critical information to get better overall performance by changing the overlay topology.

There are two possible measures that can be used to decide the shared congestion status from feedback: packet loss and delay. A packet loss based approach was proposed in [1], and a delay based approach was proposed in [2]. However, the perfect synchronization requirement for both approaches leads to limitations when they are applied as a general solution. Since both approaches assume a common source or a common destination, and a synchronization point for measuring delay data, they are susceptible to synchronization offset between probed data from two paths. Thus, they cannot be proper solutions for more general overlay topologies with multiple sources and multiple destinations.

Hence, in order to resolve the problems of the state-of-the-art techniques to detect shared congestion, a wavelet-based approach has been applied to the packet delay to detect shared congestion [3]. In this paper, the application of wavelet denoising to the multimedia network will be discussed with emphasis on the signal processing perspective. The wavelet denoising of the packet delay data allows one to filter the signal in terms of time and scale so that the filtered data is less affected by the time synchronization offset and random queueing delay fluctuations.

## II. Problem description

In Figure 1, a network topology with two paths sharing common links is illustrated. Paths $X$ from $X_{\mathrm{src}}$ to $X_{\mathrm{dst}}$ and $Y$ from $Y_{\mathrm{src}}$ to $Y_{\mathrm{dst}}$ are sharing links between $S$ and $T$. Let the one-way delay of path $X$ be $D_X$, and that of path $Y$ be $D_Y$. Each of them has two components: $d_S$, the delay from


Research sponsored by National Science Foundation ANI–0319168 and Texas Advanced Research Program 003658–0439–2001.

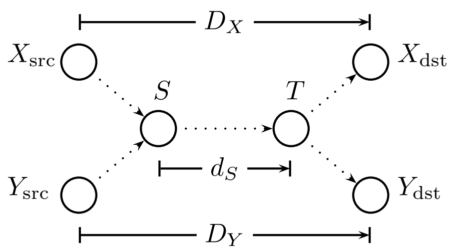


Fig. 1. A network topology with two paths sharing links

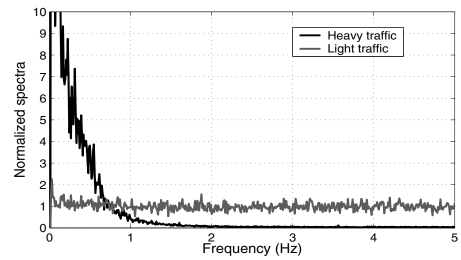


Fig. 2. Power spectra of time series with heavy traffic and light traffic

$S$ to $T$, and the remainder denoted by $d_X$ or $d_Y$.

$$\begin{aligned} D_X &= d_S + d_X \\ D_Y &= d_S + d_Y \end{aligned} \tag{1}$$

The shared congestion detection problem is to tell whether congestion occurs between $S$ and $T$. A simple but key observation is that the delay of a congested link has large fluctuations due to queueing delay changes, and if it occurs between $S$ and $T$, $D_X$ and $D_Y$ will be highly correlated. For example, if both $d_x$ and $d_y$ are constant and $d_s$ is not constant, the cross-correlation coefficient (XCOR) will be one, and if $d_s$ is constant, and $d_x$ and $d_y$ vary independently, it will be zero.

This naive cross-correlation approach is simple and seems reasonable as a measure to detect shared congestion, but it may be inaccurate due to random fluctuations of delay caused by mild congestion on non-shared links. It also requires the delay data to be probed with "perfect" synchronization. If synchronization offset is increased, the cross-correlation value drops drastically even for highly correlated delay data pair. Therefore, the interfering effects should be mitigated by wavelet denoising technique so that a shared congestion detection algorithm can be applied to more general topologies. In the following section, we discuss the nature of network delay time series data, and the application of wavelet denoising to the data.

## III. Wavelet Denoising for Detecting Shared Congestion

The frequency spectrum of the data varies in a dynamic fashion as the network traffic status changes, thus the network delay can be considered non-stationary. Figure 2 presents a pair of power spectra examples of network delay data with different network traffic status. The network traffic data are collected from the ns-2 simulator [4]. The dark line represents the average power spectrum of 30 delay data sets with heavy traffic (packet loss rate between 2% and 10%), and the light line represents the average power spectrum of 30 data sets with light traffic (no packet loss). Note that most of the power is concentrated in the low frequency band for the heavy traffic delay data, while the power spectrum of the light traffic delay data is uniformly spread over the entire frequency band. The power spectra of network delay data may vary in time between these two cases depending on the network traffic status. Therefore, any attempt to mitigate the interfering effects, such as synchronization offset and random queueing delay fluctuations, should be based on both time and frequency analysis, e.g., the wavelet transform.

The wavelet transform represents a non-stationary signal in terms of time and scale. A signal, $x(t)$, can be represented as an orthonormal expansion with wavelet basis $\psi_{i,j}(t) = 2^{-i/2}\psi(2^{-i}t - j)$ as follows [5]:

$$x(t) = \sum_{i=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} X_j^i \psi_{i,j}(t) \tag{2}$$

where the wavelet coefficients $X_j^i$ are calculated from

$$X_j^i = \int_{-\infty}^{\infty} x(t)\psi_{i,j}(t)\, dt\,. \tag{3}$$

Note that $X_j^i$ is the discrete wavelet transform of signal $x(t)$ at scale $i$ and at translation $j$.

Wavelet denoising removes noise components from the noise-corrupted signal by suppressing those wavelet coefficients of the signal which fall below a threshold. In this paper, soft thresholding is employed [6]. Wavelet denoising is based on the fact that white noise is evenly distributed over the wavelet coefficients; thus, suppressing the wavelet coefficients below the noise level can eliminate the noise from the original signal. In this paper, the interfering effects such as synchronization offset and random fluctuations are regarded as "noise," and the congestion related information is regarded as the original "signal."

However, the traditional wavelet denoising, which assumes white noise as the background noise, must be adjusted for our application to the detection of shared congestion. Since the power spectrum of the delay data with light traffic almost appears as white noise, as indicated in Figure 2, the traditional wavelet denoising technique is likely to eliminate most of this type of signal, thus leading to a meaningless result. Therefore, it is necessary to adjust the wavelet denoising approach according to the nature of the network delay data.

In recent literature concerning Internet traffic, the nature of the network traffic data is investigated in terms of time and scale distributions using the self-similarity of the data [7]–[9]. Particularly, in [8], the "energy" contained in each scale of the

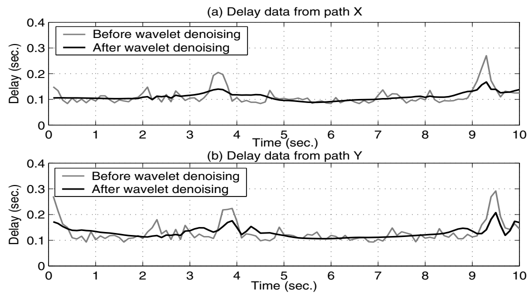


Fig. 3. Time series for shared congestion case

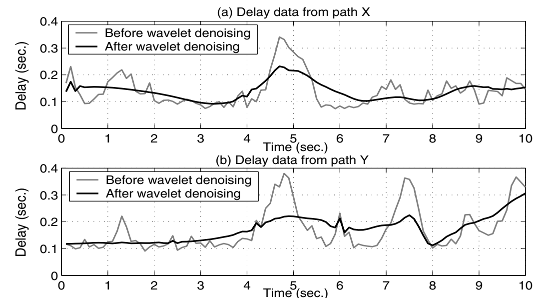


Fig. 4. Time series for independent congestion case

TABLE I
XCOR COMPARISON BEFORE AND AFTER WAVELET DENOISING

| Case | Shared congestion | Independent congestion |
|---|---|---|
| Before denoising | 0.4918 | 0.5157 |
| After denoising | 0.8027 | 0.3551 |

wavelet domain is used to interpret the network status. For the wavelet coefficients $X_j^i$ for the delay data $x(t)$, the "energy" at scale $i$ is defined as follows:

$$E_i = \frac{1}{n_i}\sum_{j=1}^{n_i}\left|X_j^i\right|^2. \tag{4}$$

where $n_i$ is the total number of coefficients at scale $i$. In [8], [10], it is shown that the self-similar traffic has a linearly increasing relationship between $\log_2(E_i)$ and scale $i$ at large values of scale $i$ (lower frequency), and queueing delay is assumed to be the main cause for this linear relationship. Hence, it is reasonable to conclude that the energy contained at large scale $i$ is associated with congestion.

Consequently, the congestion-related packet delay data can be localized at the scale of interest by choosing wavelet coefficients from those large scales. Then the remaining wavelet coefficients at smaller scales are denoised by the traditional wavelet denoising. Thus, by keeping the coefficients from large scales, we can prevent the denoised delay data from becoming meaningless when there is just slight congestion. Also, we can make our detection algorithm more robust to the interfering effects such as synchronization offset and random queueing delay fluctuations by denoising the coefficients at smaller scales.

In this paper, the wavelet coefficients from the upper half of the scales are kept for conserving congestion-related information, and the other coefficients are denoised. After denoising in the time and scale domain by the proposed wavelet denoising-based technique, the delay data from both paths are reconstructed from the selected coefficients. Then, the XCOR value between the two reconstructed delay data is calculated, and, by comparing the XCOR value to the threshold value, it can be determined whether the two paths share a common congested link.

For an illustration of how wavelet denoising makes the cross-correlation approach more robust, selected examples of wavelet denoising are provided in Figs. 3 and 4, and Table I. The original time series and the corresponding wavelet denoised time series are provided in Figs. 3 and 4, which are measured from different paths with 300 ms synchronization offset. Figure 3 is for shared congestion case, and Figure 4 is for independent congestion case. Since Figure 3 is for shared congestion case, the two time series (before wavelet denoising) in (a) and (b) should be highly correlated, but, due to the 300 ms synchronization offset, the cross-correlation value between the two time series becomes only 0.4918, as shown in Table I. However, the wavelet denoised time series have a high XCOR value, 0.8027. On the other hand, the time series in Figure 4(a) and (b) should exhibit a low XCOR value. However, due to the synchronization offset, they have a considerably high XCOR value, 0.5157, but the wavelet denoising lowers the XCOR value to 0.3551. These variations of XCOR values by the wavelet denoising are summarized in Table I.

For both cases, the XCOR values are not clearly distinguishable before the wavelet denoising, but the wavelet denoising distinguishes both cases clearly. This is because the wavelet denoising smooths the time series so that XCOR values after denoising become insensitive to synchronization offset. Furthermore, it also retains the high frequency components at those times when a strong transient peak occurs, in order to avoid the denoised time series being overly-blurred, and this explanation is verified by the comparison with a simple low-pass filter in the next section.

## IV. EXPERIMENTAL RESULTS

The performance evaluation of the wavelet-based approach is conducted by comparing the receiver operating characteristic (ROC) with 500 delay data sets. The ROC consists of a plot of probability of detection, $P_d$, vs. probability of false alarm, $P_f$. For the threshold value of cross-correlation $T_{XCOR}$, $P_d$ and $P_f$ are defined as follows.

$$\begin{aligned} P_d &= P(XCOR \geq T_{XCOR} \mid \text{shared congestion}) \\ P_f &= P(XCOR \geq T_{XCOR} \mid \text{independent congestion}) \end{aligned}$$

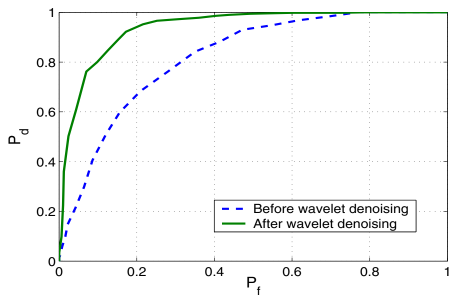


Fig. 5. Comparison of ROC before and after wavelet denoising

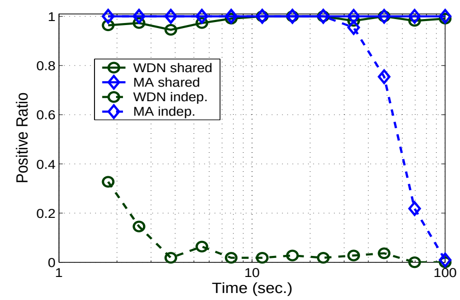


Fig. 6. Positive Ratio vs. time for wavelet denoising (WDN) and moving average (MA) for a shared congestion case, and a independent congestion case

In Figure 5, the effects of wavelet denoising on detection accuracy are compared by the ROC with 500 ms synchronization offset. The solid curve represents the ROC with the wavelet denoising, and the dashed curve represents the ROC without the wavelet denoising. The ROC area for 'after wavelet denoising' is 0.940, while the ROC area for 'before wavelet denoising' is 0.821. This implies that the wavelet-based detection algorithm exhibits almost ideal detection performance even with 500 ms synchronization offset, whereas the detection performance without the wavelet denoising is unreliable.

Wavelet denoising works as a low-pass filter with severe congestion. However, simple low-pass filtering may not work with light traffic. Without congestion, for instance, near-zero cross correlation is expected since the delay signal will be dominated by random queue behaviors. The wavelet denoising preserves those dominant high-frequency components, resulting in a low XCOR value. On the other hand, low-pass filtering may over-blur the delay signal.

To verify the effect of wavelet denoising, the convergence speed of wavelet denoising and a simple low-pass filter, i.e., moving average, is compared. To compare the accuracy of each technique, we define a new metric as follows.

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

For a fair comparison, the span of the moving average (1.1 sec.) was selected so that its effect on the synchronization offset is similar to that of wavelet denoising shown in Figure 5. In Figure 6, we plotted the Positive Ratio for wavelet denoising (WDN) and moving average (MA) techniques over 100 experiments for shared congestion and independent congestion as time progresses.

The moving average increases the correlation. Thus, the positive ratio of moving average with shared congestion is very close to one from the beginning. However, the positive ratio of moving average with independent congestion is incorrectly near one for 30 seconds before it begins to drop to its correct value of zero. On the other hand, the modified wavelet denoising case is always close to one for the shared congestion case, and quickly (after 3–4 seconds) approaches the correct value of zero for the independent congestion case.

## V. Conclusion

In this paper, we discussed the feasibility of wavelet denoising to detect shared congestion in overlay networks, and suggested a modification of wavelet denoising for the network delay data. From Figs. 5 and 6, one can conclude that the modified wavelet denoising technique is robust to synchronization offset and queueing delay fluctuations, and exhibits faster convergence. In addition, the accuracy of the detection is enhanced via the wavelet denoising. This successful result is due to the ability of the proposed wavelet denoising to reduce the interfering effects such as synchronization offset and random queueing delay fluctuations. We believe that many multimedia overlay systems can utilize network resources more effectively by the use of the proposed wavelet denoising technique.

## References


[1] K. Harfoush, A. Bestavros, and J. Byers, "Robust identification of shared losses using end-to-end unicast probe," in *Proceedings of the 8th IEEE International Conference on Network Protocols*, Nov. 2004.

[2] D. Rubenstein, J. Kurose, and D. Towsley, "Detecting shared congestion of flows via end-to-end measurement," *IEEE/ACM Transactions on Networking*, vol. 10, pp. 381–395, June 2002.

[3] Min Sik Kim, Taekhyun Kim, YongJune Shin, Simon S. Lam, and Edward J. Powers, "A wavelet-based approach to detect shared congestion," in *Proceedings of ACM SIGCOMM 2004*, Aug. 2004.

[4] S. McCanne and S. Floyd., "ns network simulater." [Online]. Available: http://www.isi.edu/nsnam/ns/

[5] I. Daubechies, "The wavelet transform, time-frequency localization and signal analysis," *IEEE Transactions on Information Theory*, vol. 36, no. 5, pp. 961–1005, Sept. 1990.

[6] D. L. Donoho, "De-noising by soft-thresholding," *IEEE Transactions on Information Theory*, vol. 41, no. 3, pp. 613–627, May 1995.

[7] W. E. Leland, M. S. Taqqu, W. Willinger, and D. B. Wilson, "On the self-similar nature of ethernet traffic (extended version)," *IEEE/ACM Transactions on Networking*, vol. 2, pp. 1–15, Feb. 1994.

[8] P. Arby and D. Veitch, "Wavelet analysis of long-range dependent traffic," *IEEE Transactions on Information Theory*, vol. 44, pp. 2–15, 1998.

[9] M. S. Borella and I. Sidhu, "Self-similarity of Internet packet delay," in *IEEE International Conference on Networking*, Jan. 1997, pp. 513–517.

[10] Q. Li and D. L. Mills, "On the long-range dependence of packet round trip delays in Internet," in *IEEE International Conference on Communications*, 1998, pp. 1185–1191.