

# Accurate Shared Bottleneck Detection Based On SVD and Outliers Detection

Muhammad Murtaza Yousaf, Michael Welzl, and Bülent Yener

**Abstract**—We present a new mechanism for detecting shared bottlenecks between end-to-end paths in a network. Our mechanism, which only needs one-way delays from endpoints as an input, is based on the well known linear algebraic approach SVD (Singular Value Decomposition). Clusters of flows which share a bottleneck are extracted from SVD results by applying an outlier detection method. Simulations with varying topologies and different network conditions show the high accuracy of our technique.

**Index Terms**—Shared bottleneck detection, network monitoring, passive measurement, singular value decomposition, outliers detection.

## I. INTRODUCTION

**E**ND-TO-END flows in packet based networks adversely influence each other when they traverse a network link which is a single point of congestion. We call such a link a “shared bottleneck”. Given the large number of distributed applications that are now used on the Internet, where coordination between such end-to-end flows is quite feasible, we believe that the detection of shared bottlenecks in the network is one of the key measurement efforts that can be undertaken, as there are numerous possibilities for using such knowledge.

Consider, for example, a transfer delay prediction tool (which is a common element of computational Grids) which tells two senders A and B that the time to transfer files to receivers C and D will be four and five minutes, respectively, but fails to inform them that the A-C and B-D path share a bottleneck. Then, both systems would perhaps start to send the files at the same time, and the predicted transfer delay would clearly be wrong because of adverse interactions between the flows at the shared bottleneck. In a Grid, it is not uncommon for a single entity (or person) to be in control of the two senders A and B — thus, with knowledge about shared bottlenecks, this problem might have been prevented by accordingly scheduling the time of these transfers (this is what schedulers in Grid workflow systems do).

As another example, one could imagine an overlay network where a file is sent from a single source to a single destination via multiple paths across the overlay in order to speed up the transfer time. Such a scheme would probably only yield a noticeable benefit if the overlay paths that are used do not include a shared bottleneck in the underlay. Similarly,

M. M. Yousaf and M. Welzl are with the Institute of Computer Science, University of Innsbruck, Austria (e-mail: murtaza.yousaf@uibk.ca.at; michael.welzl@uibk.ac.at).

B. Yener is with the Rensselaer Polytechnic Institute, Lally 310, 110 Eighth Street, Troy, NY 12180-3590, USA (e-mail: yener@cs.rpi.edu).

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

in peer-to-peer systems, knowledge about shared bottlenecks could theoretically be used for appropriately rearranging the way documents are placed in order to exploit multiple paths or avoid adverse interactions among simultaneously active downloaders.

To the best of our knowledge, despite the existence of some measurement studies on shared bottleneck detection (cf. [1], [2] and [3]), such endeavors are hardly undertaken. This may be due to problems with the existing measurement methods, such as the amount of data that they need to transfer, or the time that a measurement would take. Most notably, the accuracy exhibited in recent studies is limited — this may be a reason for designers of new network mechanisms to refrain from using such measurement results in practice.

In this paper, we address this issue with a method that only operates on end-to-end measurements (forward delays), is easy to calculate and provides extraordinarily accurate results. Our scheme is based on a combination of the well known technique SVD (Singular Value Decomposition) and outlier detection.

After a brief survey of related work in the next section, we will provide the mathematical foundations of our mechanism in section 3. The mechanism itself will be explained in section 4, and results from simulations will be presented in section 5. Section 6 concludes with an outlook on future work.

## II. RELATED WORK

The awareness of the complete network topology and particularly the knowledge of congested network links which are shared by multiple flows can be of great use to many applications. A great deal of work has already been done in this area which directly or indirectly relates to our study. In what follows, we briefly discuss some of the related efforts.

Two techniques have been proposed in [1] which are mainly based on correlation: one is based on the forward delays of the packets and the other relies on packet loss. The decision of whether two flows share the same bottleneck or not is made by a comparison test. Such a test would compare the cross measure (cross correlation) and auto measure (auto correlation), so the decision is not only taken on the basis of a single correlation function.

In the delay based technique the cross measure is computed by merging the delay sequences of the two flows which are to be tested. From this merged sequence only those pairs of samples are picked in which the first sample of a pair is from one sequence and the second sample of pair is from another sequence. Then, from all such samples, the cross measure is calculated between the samples of the first and the second sequence. The auto measure is calculated as the correlation

coefficient between the delays of adjacent packets of a single flow. After the computation of the cross measure and the auto measure, the comparison test declares that the two paths share a bottleneck if the auto measure is less than the cross measure.

In the loss based technique, the cross measure and the auto measure are actually the conditional probabilities that a packet is lost given that the following packet is also lost. Both techniques work in the topologies where the senders or receivers are co-located. It is also shown that the delay based technique attains accuracy quite quickly as compared to the loss based technique.

The entropy based technique presented in [4] works in topologies where some sources are connected to a single destination with the assumption that each path holds only one bottleneck. Passive measurements are used in this technique to measure the inter-arrival times at the destination. A cluster of flows is formed and the average entropy is calculated for every flow with the assumption that the flow belongs to this cluster. After that the flow is put into the cluster with the minimum average entropy. The clustering correctness of this technique only gives acceptable results in case of low cross traffic and does not scale well with heavy cross traffic.

We consider FlowMate [3] as one of the most relevant works for us. It uses the comparison test for the delay based technique of [1], and works on the online clustering of flows by comparing each new flow with a representative flow of every cluster. Initially, one flow is designated in each cluster and later every new flow is compared with only one representative of each cluster. Another simple clustering mechanism is also discussed in [3] where representatives of each cluster can be switched dynamically. The accuracy of FlowMate however becomes a problem in case of short lived HTTP connections because it does not get enough samples.

One of the latest works in this regard is by Kim et al. [5], where a clustering approach which uses multidimensional indexing is presented. This work is based on the previous work presented in [2], where the cross correlation was used to detect whether two paths share a bottleneck. This cross correlation is computed between the delay sequences of two paths after applying wavelet denoising on the original delay sequences.

One thing which is common in all above mentioned approaches is that all of them operate in a pairwise fashion. All techniques consider two paths to detect if flows share a bottleneck whereas our mechanism works for any number of flows at the same time, resulting in clusters of paths which share a bottleneck.

A loss based technique which expects a common sender to send packet pairs has also been proposed [6]. If the probability of the second packet being lost exceeds a particular threshold, it is assumed that the two paths share a point of congestion. This technique requires DropTail queuing and can suffer from all the problems which exist for packet pair based measurements.

### III. MULTIVARIATE ANALYSIS

Multivariate analysis has most of its applications in image processing and pattern recognition but the nature of our problem makes multivariate analysis an ideal tool in solving it. We

decided to apply multivariate analysis because we believe that there are many factors which are beyond the scope of a simple correlation method. Multivariate analysis is highly suitable in such scenarios because we want to check the principal factors which mean the delays caused by the congested link.

We will see that SVD provides us with highly accurate results in terms of identifying the flows which share a bottleneck. SVD is one out of many multivariate analysis techniques in the literature, and others should eventually also be considered in terms of efficiency and reliability.

We construct a matrix  $A$  in which the rows are the flows and the columns are time slots. An entry in  $A$  is the amount of (average) delay measured during experiments. We will analyze this data matrix  $A$  using SVD.

#### A. SVD and Spectral Analysis

We briefly review the singular value decomposition of matrices; we will use some of its properties in our discussion in section IV-B.

According to [7] any  $m \times n$  matrix  $A$  can be expressed as

$$A = \sum_{t=1}^r \sigma_t(A) u^{(t)} v^{(t)T},$$

where  $r$  is the rank of  $A$ ,  $\sigma_1(A) \geq \sigma_2(A) \geq \dots \geq \sigma_r(A) > 0$  are its singular values and  $u^{(t)} \in \mathcal{R}^m, v^{(t)} \in \mathcal{R}^n, t = 1, \dots, r$  are its left and right singular vectors respectively. The  $u^{(t)}$ 's and the  $v^{(t)}$ 's are orthonormal sets of vectors; namely,  $u^{(i)T} u^{(j)}$  is one if  $i = j$  and zero otherwise. We also remind the reader that

$$\|A\|_F^2 = \sum_{i,j} A_{ij}^2 = \sum_{i=1}^r \sigma_i^2(A),$$

$$\|A\|_2 = \max_{x \in \mathcal{R}^n: \|x\|=1} \|Ax\| = \max_{x \in \mathcal{R}^m: \|x\|=1} \|x^T A\| = \sigma_1(A).$$

In matrix notation, SVD is defined as  $A = U\Sigma V^T$  where  $U$  and  $V$  are orthogonal (thus  $U^T U = I$  and  $V^T V = I$ ) matrices of dimensions  $m \times r$  and  $n \times r$  respectively, containing the left and right singular vectors of  $A$ .  $\Sigma = \mathbf{diag}(\sigma_1(A), \dots, \sigma_r(A))$  is an  $r \times r$  diagonal matrix containing the singular values of  $A$ .

One of numerous applications of SVD is in recovering the structure of matrices and in noise removal in a variety of settings. The underlying idea is very simple: since  $A = \sum_{t=1}^r \sigma_t(A) u^{(t)} v^{(t)T}$ , we can create approximations to  $A$  by keeping only the top  $k$  "principal components" (i.e. the top  $k$   $\sigma_t(A) u^{(t)} v^{(t)T}$ ) for various values of  $k$ . Essentially, discarding the "smallest" principal components (the ones corresponding to the smallest singular values), results in a small loss in accuracy, and we might justifiably consider these components as "noise". This procedure is commonly referred to as "Principal Component Analysis" and the following theorem (usually attributed to Eckart and Young [7]) quantifies the loss of accuracy incurred by keeping only the top  $k$  components for various values of  $k$ .

*Theorem 1:* Let  $A_k = \sum_{t=1}^k \sigma_t(A) u^{(t)} v^{(t)T}$  (for any  $1 \leq k \leq r$ ).  $A_k$  is the "best" rank  $k$  approximation to  $A$  with respect

to the 2-norm and the Frobenius norm; namely, for any matrix  $D$  of rank at most  $k$ ,

$$\|A - A_k\|_2 \leq \|A - D\|_2 \quad \text{and} \quad \|A - A_k\|_F \leq \|A - D\|_F.$$

Also,

$$\|A - A_k\|_F^2 = \sum_{t=k+1}^r \sigma_t^2(A) \quad \text{and} \quad \|A - A_k\|_2 = \sigma_{k+1}(A).$$

We say that a matrix  $A$  has a “good” rank  $k$  approximation if the 2-norm and the Frobenius norm of  $A - A_k$  is small; for a detailed treatment of Singular Value Decomposition see [7].

### B. Intuition

The main intuition behind using SVD based approach was obtained from the work of Ahmet et. al. [8] where the authors used SVD based approach to discover hidden communities and communication patterns in a chatroom. A tool was built for the data collection of individual chatters, and then an SVD based algorithm was applied to locate groups of chatters communicating with each other. The authors used the number of messages exchanged during a time window as data and found the groups of chatters communicating with each other on the basis of a similar number of messages exchanged. We mapped this to our problem by comparing different nodes in the network with the chatters in the chatroom and considering forward delays along a path as data entries during a time window, assuming that this should provide us with an association among the paths with similar delay patterns during certain periods of time. As the end-to-end delay along a network path is mainly influenced by the bottleneck, we assumed that SVD must cluster the paths with similar delay patterns during a certain period of time or in other words, that share the same bottleneck.

## IV. SVD BASED SHARED BOTTLENECK DETECTION

A bottleneck is a network link/router where the packets of a flow experience excessive delays or they are even lost. Most of the techniques proposed till now to detect shared bottlenecks are based on packet delays or packet loss. It has been shown that delay based techniques quickly converge towards accurate results [1]; hence we decided to rely on the same metric in our mechanism, where one-way delays along a path are used to detect shared bottlenecks. The whole approach is summarized in Algorithm 1. We can divide it into three major steps:

- 1) Measurement of one-way delays along all the paths
- 2) Applying SVD on the matrix delay measurements of all the paths and getting a projection
- 3) Grouping paths that share a bottleneck

### A. Delay Measurements

First of all we measure the one-way delays for all the paths by initiating a flow on every path. As our algorithm requires full knowledge about these measurements, we assume that they are then then collected in a single place — e.g., at the server which will be queried by users for measurement results. There might be interesting ways of building a more scalable

distributed measurement infrastructure, but at this point we stick with the assumption of a single server for the sake of simplicity. We group samples of delay measurements on the basis of time windows and pick the average of a time window as a representative delay measurement for that particular time window. We have worked with different time windows and observed no significant change in our results; this is explained in the section V-B3. After checking with different window sizes, we decided to use 0.3 seconds as an optimal value for realistic network conditions.

Clearly, this value is a trade-off between accuracy (the smaller the window, the more information we have) and overhead (we need to collect the necessary data at the network node that will carry out the calculation). We can also reduce or expand the time window size depending on the delay samples available, for instance in case of short lived flows we can reduce it and in case of long lasting flows we can expand it, e.g., if we need to work in a fully passive environment. At the end we form a matrix  $D$  of forward delays of all flows as shown in equation (1) where each row represents the delays experienced by a single flow. Each entry  $d_{ij}$  is an average delay measurement of the  $i^{\text{th}}$  path for the  $j^{\text{th}}$  time window.

$$D = \begin{bmatrix} d_{11} & d_{12} \cdots & d_{1n} \\ d_{21} & d_{22} \cdots & d_{2n} \\ \vdots & & \\ d_{m1} & d_{m2} \cdots & d_{mn} \end{bmatrix} \quad (1)$$

### B. Applying SVD and Projecting

Having the matrix  $D_{m \times n}$  of delay measurements, we apply SVD on  $D_{m \times n}$  to get singular values of our delay measurements for all the paths. As a result we get the left singular matrix  $U_{m \times r}$ , the diagonal matrix  $\Sigma_{r \times r}$ , and the right singular matrix  $V_{n \times r}$ . Here, the matrix  $U_{m \times r}$  represents the association among the paths on the basis of similar delays caused by common bottlenecks, and the matrix  $\Sigma_{r \times r}$  represents the degree of this association. Then, on the basis of rank reduction of the diagonal matrix  $\Sigma_{r \times r}$ , we project the matrix  $U_{m \times r}$  on the  $c$  column(s) (for any  $1 \leq c \leq r$ ) of the diagonal matrix to look for  $c$  dimensions. The value of  $c$  can be decided by looking at the explained variance of the elements of the diagonal matrix. For instance, if the first element of the diagonal matrix is explaining the data 60% then we may need to take the second or third dimensions. If the first one explains 95% of the data we say there is *rank reduction* and project on the first column of the diagonal matrix to take only the first dimension.

Usually, if the paths under investigation share a bottleneck then the diagonal matrix always shows a good *rank reduction* and we need to project only on the first column of diagonal matrix as shown in equation 2.

$$X_{m \times 1} = U_{m \times r} \times \Sigma_{r \times r}(:, 1) \quad (2)$$

In order to project matrix  $U_{m \times r}$  on the  $c^{\text{th}}$  column of the diagonal matrix, we multiply each row of  $U_{m \times r}$  with only the  $c^{\text{th}}$  column of the diagonal matrix. This results in

---

**Algorithm 1** Detection of paths that share a bottleneck
 

---

**Input:**  $D_{m \times n}$  {A matrix of  $n$  successive delay measurements for  $m$  network paths}  
**Output:**  $P = \{P_1, P_2, \dots, P_{j \leq m}\}$  Where  $P_i$  is set of paths that share the bottleneck  
 $[U, \Sigma, V] \leftarrow \text{svd}(D)$   
 $c \leftarrow \text{columnToProject}$   
**for**  $i = 1$  to  $m$  **do**  
      $x_i \leftarrow u_{ic} \times \sigma_{cc}$   
**end for**  
 $X = \{x_1, x_2, \dots, x_m\}$  {A set of projection values corresponding to each path}  
 $X \leftarrow \text{sort}(X)$   
**for**  $i = 1$  to  $m-1$  **do**  
      $z_i \leftarrow \frac{x_i - x_{i+1}}{x_{i+1}}$   
**end for**  
 $Z = \{z_1, z_2, \dots, z_m\}$  {A set of *rate of change* of projection values corresponding to each path}  
 $\bar{z} \leftarrow \text{mean}(Z)$   
 $s \leftarrow \text{std}(Z)$   
 $\tau \leftarrow \bar{z} + 0.2 \times s$  {Threshold to identify group of paths that share a bottleneck }  
 $j \leftarrow 1$   
 $P_j \leftarrow \phi$  { $j^{\text{th}}$  group of paths that share a bottleneck}  
**for**  $i = 1$  to  $m$  **do**  
     **if**  $z_i > \tau$  **then**  
          $P \leftarrow P_j$   
          $j \leftarrow j + 1$   
          $P_j \leftarrow \phi$   
     **end if**  
      $P_j \leftarrow p_i$  {  $p_i$  is path corresponding to  $z_i$  }  
**end for**  
**return**  $P$

---

a matrix  $X_{m \times 1}$ , as shown in equation 3, that has one entry corresponding to each path and these values are further used to group flows on the basis of shared bottlenecks.

$$X = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{bmatrix} = \begin{bmatrix} u_{11} & u_{12} \cdots & u_{1r} \\ u_{21} & u_{22} \cdots & u_{2r} \\ \vdots & & \\ u_{r1} & u_{r2} \cdots & u_{mr} \end{bmatrix} \begin{bmatrix} \sigma_{11} & 0 \cdots & 0 \\ 0 & \sigma_{22} \cdots & 0 \\ \vdots & & \\ 0 & 0 \cdots & \sigma_{rr} \end{bmatrix} \quad (3)$$

### C. Grouping Flows

As already stated, the matrix  $X_{m \times 1}$  contains one value for each path. Here the values for paths that share a bottleneck are similar, and they differ significantly from the values of other paths. A simple plot of such values will look as in figure 1a. If we sort all the entries of matrix  $X_{m \times 1}$  in descending order and plot the sorted data, then even at this step we can easily visualize some clusters as exhibited for the same data in figure 1b, which shows our result of one of the experiments with 25 paths and 3 bottlenecks. If we visualize the data as a graph we can observe some clear grouping on the basis of

stable slopes and sudden drops in the slope indicating the start of a new cluster.

We find the *rate of change* (*roc*) between two successive values of the sorted list by using

$$roc = \frac{x_i - x_{i+1}}{x_{i+1}} \quad (4)$$

The *rate of change* could be calculated by dividing the successive differences with the denominator being either  $x_i$  or  $x_{i+1}$ . As the values have already been sorted in descending order, it is obvious that this series exposes the break points better if we choose  $x_{i+1}$  as denominator as compared to  $x_i$  which is shown in figure 1c. Once we have this series of *rate of change*, the next step is to detect the break points which becomes a simple problem of detecting outliers.

In order to detect outliers we used the threshold  $\tau$  defined as:

$$\tau = \bar{z} + 0.2 \times s \quad (5)$$

Where  $\bar{z}$  is the mean *rate of change* and  $s$  is its standard deviation. The coefficient value of 0.2 was selected after a detailed study of data obtained in simulations. As we will see in section V, our method works very well with this value in a wide range of network conditions and scenarios.

### D. Time complexity of the whole Algorithm

It is important to have a computable time complexity for situations where a mechanism needs to be applied in real time. In this section we discuss the time complexity of our algorithm and we show that it is feasible to use online. It has three major parts. During the first part, we extract singular values from delay measurements, then in the second phase we find the projection of left singular vectors over the first singular value, and finally in the third and last phase we identify clusters of paths that share a bottleneck.

First of all let us discuss the time complexity of SVD which is  $O(mn^2)$  for a matrix of order  $m \times n$  [7]. Here, the dominant part is the number of columns of a matrix. In our technique, the number of columns or  $n$  is the number of delay measurements used to detect shared bottlenecks. In the section V we will show that very few measurements (1 measurement is the average of all samples taken in 300 milliseconds) are enough to obtain an accurate result. The time complexity of  $O(mn^2)$  is not expensive at all when we have an upper bound on  $n$ , so if we consider  $n$  as a constant then the complexity becomes  $O(m)$ , which means that it depends only on the number of paths under investigation.

In an online version, we may need to update our results in a continuous fashion because the bottleneck link may vary depending on the load of network traffic. Solving SVD for one matrix of delay measurements allows for an incremental, efficient solution for letting the matrix slide over in time (i.e., throwing away the oldest measurements and bringing in the most recent). In such situations, the time complexity of SVD is further reduced to  $O(mk)$  [9] where  $k$  is the size of the data set.

The second phase involves the projection which mainly multiplies the first element of each row of matrix  $U$  with the

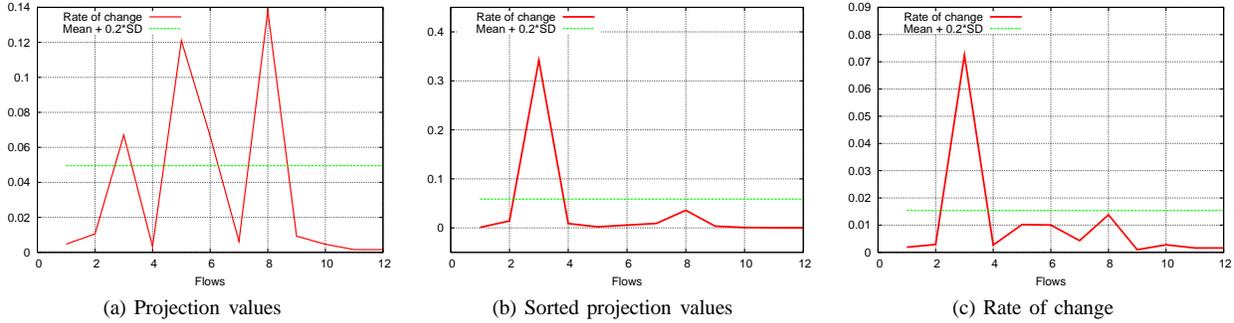


Fig. 1. Results in intermediate and final form

first element of the first column of matrix  $\Sigma$ . This procedure takes  $m$  multiplications which means a time complexity of  $O(m)$ .

We get  $m$  values corresponding to each path as a result of second stage. In the final part of our procedure we first sort these values, which may take  $O(m \log m)$  steps (e.g. if quicksort is used). At the end, for grouping paths which share a common bottleneck, we need to compare the sorted list with a threshold. This can be done in  $O(m)$  steps because it just involves  $m$  comparisons. We perform all above steps in serial so the most dominant part defines the time complexity of our whole algorithm. The total time complexity is therefore  $O(m \log m)$  where  $m$  is the number of paths to be inspected for common bottlenecks. This shows that it is practical to use this approach in real time.

### V. PERFORMANCE EVALUATION

In this section, we present the results of using our mechanism in simulations. We used the ns-2 [10] network simulator to examine different scenarios. We have applied the same simulation topologies and parameters which were investigated in [3] because we consider it as one of the most relevant works to us. One of the two topologies used in [3] is simple and symmetric, while the other one is more complex and asymmetric.

The symmetric topology is shown in figure 2a. It contains three different bottlenecks on three different branches of paths. Each branch has three different destinations for a common source. All the links capacities are 10 Mbps except for the bottleneck links of the first two branches. The bottleneck links at the first two branches have a capacity of 1.5 Mbps and 3.0 Mbps respectively, whereas the third bottleneck is created by heavy cross traffic from a number of constant bit rate (CBR) flows. Background traffic is generated in both directions from source to destinations by a number of multiplexed Pareto flows. The rest of the simulation parameters are summarized in table I.

The asymmetric topology is shown in figure 2b. It also has three different branches with a bottleneck at each branch. The mechanism for creating the bottlenecks is the same as for the symmetric topology. It has a varying number of destinations, and for each destination, it has a varying number of links for the end-to-end path. Simulation parameters are given in table I

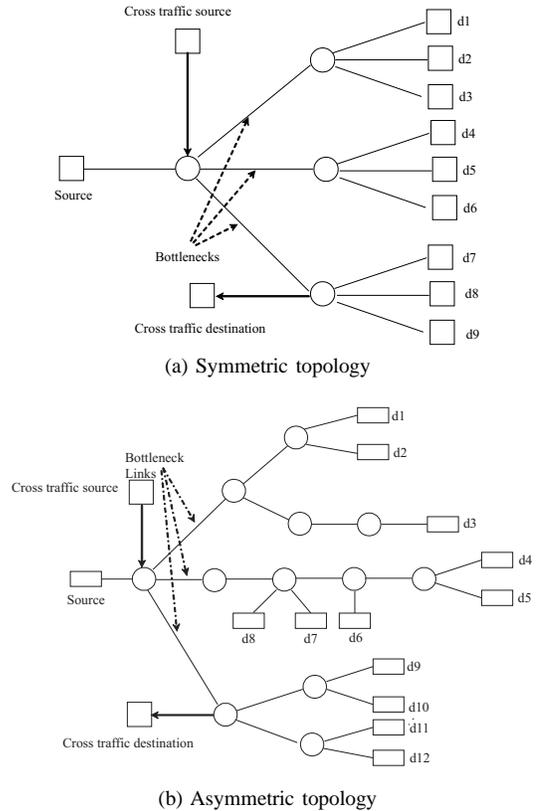


Fig. 2. Symmetric and asymmetric topologies

except for the number of foreground TCP flows which is 12-36.

#### A. Performance Metrics

Three different performance metrics which can be considered for measuring the accuracy have been used in [2]:

- **False positive** — a flow does not share a bottleneck but it is considered as being part of a group that shares a bottleneck.
- **False negative** — a flow shares a bottleneck with a group but is not identified as part of the group.
- **Correct clustering** — a flow is in the same group as flows with which it shares a bottleneck.

The Accuracy Index ( $AI$ ) presented in [3] takes all of the above performance metrics into account and measures the

TABLE I  
SIMULATION PARAMETERS

Foreground flows	9-27 TCP flows
Background traffic	256 Kbps Pareto flows (source to destinations)
Reverse traffic	64 Kbps Pareto flows (destinations to source)
Cross traffic	64 Kbps CBR flows (24)
Queue size	250 packets (In one experiment 50, 100, 150, .... 250)
Drop policy	DropTail (One experiment with RED and RED + DropTail)
Link Delays	Randomly assigned (From 1ms to 20ms)

accuracy of clustering in a unified manner. If there are  $N$  paths,  $P_c$  is the set of correct clusters,  $P_o$  is the set of clusters as a result,  $n_{fs}$  is the number of flows erroneously included in a resulting cluster, and  $s_j$  is the number of sub-clusters of a correct cluster that was split into  $s_j$  sub-clusters in  $P_o$ , then  $AI$  is calculated as:

$$Accuracy\ Index\ (AI) = 1 - \frac{\sum_{i=1}^{|P_o|} (n_{fs})_i}{N} - \frac{\sum_{j=1}^{|P_c|} (s_j - 1)}{N}$$

## B. Results

In this section, we discuss the results of our experiments in different ways to highlight the accuracy of our SVD-based method for detecting shared bottlenecks.

1) *Symmetric Topology*: To illustrate our whole mechanism, we first used the symmetric topology shown in figure 2a with 9 TCP flows (one for each destination). In this experiment, flows 1-3, 4-6 and 7-9 each share a bottleneck. For better understanding, in the first row of table II we show the values of the diagonal matrix  $\Sigma$  which exhibits a nice rank reduction, whereas in the second row the values of matrix  $V$  are shown; these are the results after projection of the left singular matrix  $U$  on the first component of the diagonal matrix  $\Sigma$  (as explained in section IV-B). Just from a brief look at the values of  $V$  in table II, we can visualize a grouping of flows 1-3, 4-6, and 7-9.

As the next and final step, we calculate the rate of change for the values of matrix  $V$  as explained in section IV-C and find the outliers to detect the group of flows which share a bottleneck. The results are shown together with the absolute values of  $V$  in figure 3a; they are accurate according to the simulation scenario.

We conducted the same experiment with 18 and 27 foreground TCP flows to check the effect of multiple TCP flows over the same bottleneck. In these experiments we started 2 and 3 TCP flows for the same destination, using a random start time within the first second of the simulation in order to minimize phase effects [11]. In both experiments, we obtained 100% accurate shared bottleneck recognition (i.e. Accuracy Index  $AI = 1$ ). The results are depicted in figure 3.

2) *Asymmetric Topology*: We performed experiments with 12, 24, and 36 flows with random start time within the first second of the simulation using the asymmetric topology and found  $AI$  to be always 1. The results are shown in figure 4.

3) *Window Size*: We used the asymmetric topology with 12 foreground TCP flows [representing 12 paths]. In this experiment we wanted to check the impact of the window size on the accuracy of our method. By window size we mean the

time interval for which we combine all delay measurements and choose one representative which is the average of that interval. A window size of 0.3 seconds is working fine, in what follows, we justify the use of this value. We ran the simulation for 60 seconds and obtained results after setting different window sizes. It is possible that, if we keep on reducing the size of the window, we will reach a point when some windows may not have even a single measurement due to congestion, meaning we will not have even a single delay measurement for that particular window. This must be avoided because the SVD may result in inaccurate results if we use some blanks or imaginary values in the matrix of delay measurements. At the same time, we should not make the window size too large because this could blur the information about the correct bottleneck and the time series of delay measurements of different flows may not show any association.

During this experiment we varied the window size from 0.1 seconds to 1.0 seconds, increasing it by 0.1 seconds. The resulting accuracy index was 1 after a window size of 0.2 onwards, as shown in figure 6a. Figure 6b shows very stable slopes with significant rate of change  $roc$  for projected values for all the windows sizes we investigated. This makes it much easier to identify cluster of flows that share a bottleneck.

As already mentioned, we started from a very small window size of 0.1. In this case there were many windows for which we could not have even a single delay measurement. We placed -1 in the empty windows because SVD does not work with empty cells in the matrix. We show the percentage of negative values (i.e. empty windows) against the window size in figure 6c. We notice that the percentage of empty cells for a window size of 0.2 is 2.41%, but even then we got accurate results as depicted in figure 6a. This means that our method is strong enough to cope with some percentage of missing data less than or equal to 2%.

On the basis of the above observations we propose a general principle to choose a window size for any environment such that it may not yield a percentage of empty cells of more than 2%. In the rest of our experiments we used a window size of 0.3 which worked fine with all the dynamic network conditions and environments that we simulated. Another strength of our approach revealed in the above experiments was the amount of data required to converge to accurate results. While increasing the window size of 30 seconds, which means that we had only 2 average measurements per flow. In such situation, we had a delay matrix of  $12 \times 2$  and it resulted in only 2 values in the diagonal matrix  $\Sigma$ , but even with such a small number of values in the diagonal matrix we had perfect results after the projection step.

TABLE II  
THE VALUES OF THE DIAGONAL MATRIX AND THE FINAL RESULT AFTER PROJECTION

Flows	1	2	3	4	5	6	7	8	9
$\Sigma$	9.9184	1.1069	0.3380	0.2304	0.1663	0.0645	0.0329	0.0308	0.0251
Projection Result	-5.0970	-5.1072	-5.1147	-2.4315	-2.4323	-2.4201	-0.9030	-0.8959	-0.9206

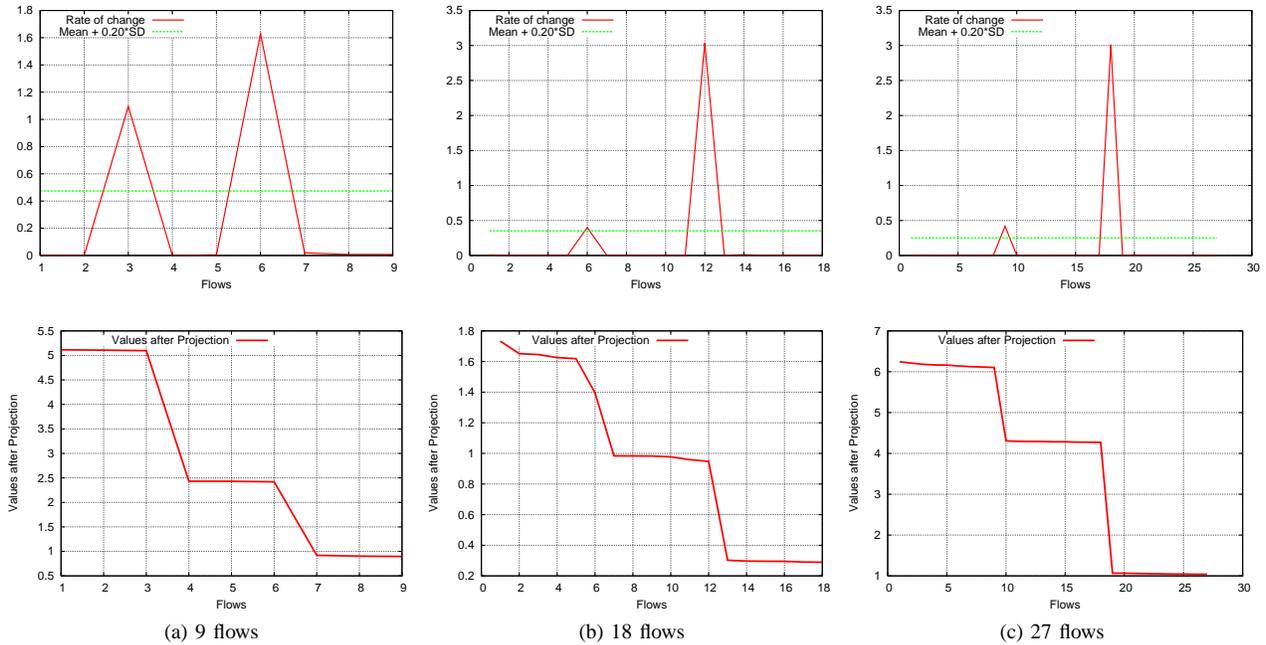


Fig. 3. Results of symmetric topology with varying number of flows

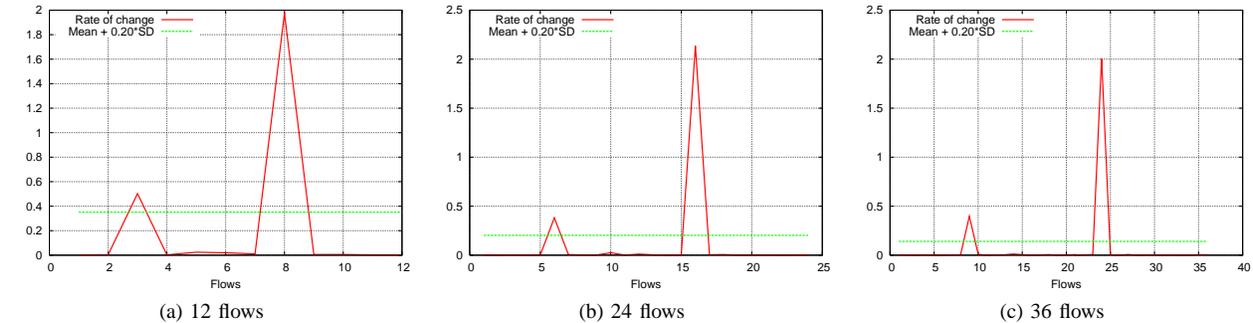


Fig. 4. Results of asymmetric topology with varying number of flows

4) *Drop Policy*: We observed some interesting results in simulations using the asymmetric topology with different drop policies like RED (according to the recommendation in [12], we allowed the algorithm implemented in ns-2 to automatically configure the parameters (“q\_weight -1”, “thresh”, and “maxthresh”, for automatic configuration)) and DropTail. DropTail performed well and showed a very promising behavior as shown in figure 4. However in the case of Random Early Detection (RED), while adopting our usual procedure we found a small margin of inaccuracy. This inaccuracy occurred because we could not reach the desired rank reduction in the result of SVD. Thus, we projected our data on 2nd and 3rd components, and consequently the results became more accurate as shown in figure 5.

It is not surprising that our mechanism has problems with RED, which improves fairness and reduces phase effects by randomly dropping packets [13] — any technique that tries to detect dependencies between flows will naturally suffer from a mechanism that is designed to work against these dependencies. Nevertheless, some dependencies may still remain; thus, although DropTail is commonly known to be the most common drop policy used in routers, we feel that we should work on the automatization of our mechanism to work with all possible drop policies by introducing some additional decision making during the projection phase. This is planned as future work.

5) *Buffer Size*: Our additional tests included scenarios with different buffer sizes from 50 to 500 packets. The results showed no significant difference. With DropTail, the AI was

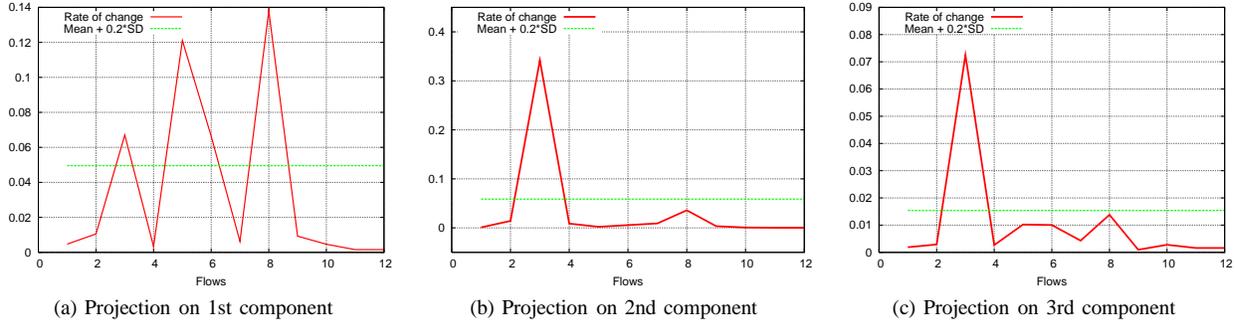


Fig. 5. Results of asymmetric topology with RED as drop policy and projecting on different components of diagonal matrix

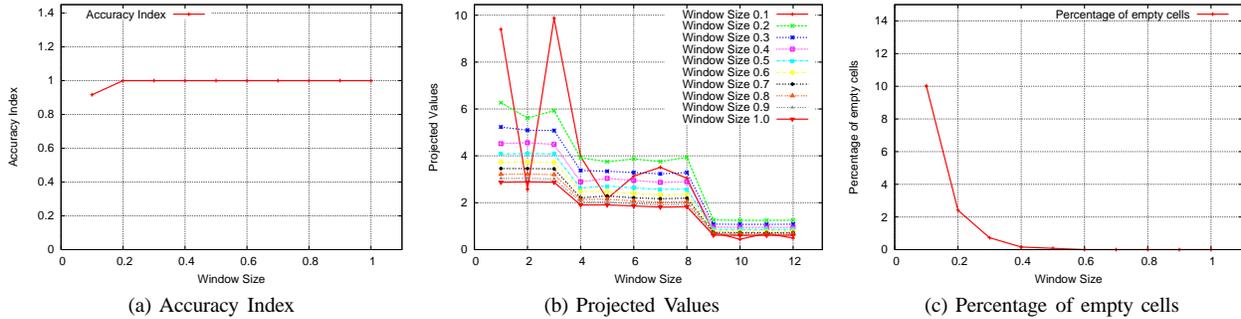


Fig. 6. Results of window size experiments

always 1. We also computed the *AI* after every 5 seconds for all above scenarios, starting from second 1 of the simulation till its end. We found that the *AI* value was 1 throughout the simulations.

6) *Without Common End Points*: In our experiments so far, we have used topologies with a common source. It has also been the major assumption in related work to have a single common source or destination. Here, we evaluate our technique for scenarios with multiple sources and destinations. The topology used for these experiments is shown in figure 7. We have created it by considering the mirror image of all three branches of the asymmetric topology shown in figure 2b. The center of the topology represents an overprovisioned wide-area network backbone, whereas bottlenecks are placed at the access points of local networks, which resembles the most common real life case of Internet connectivity. Two of the three bottlenecks are limited because of their small physical capacity (1.5 Mbps and 3.0 Mbps, respectively) whereas the third bottleneck is caused by heavy cross traffic. The parameters and details for cross traffic and background (forward and reverse) traffic are the same as with previous topologies — these parameters are mentioned in table I except for the non-bottleneck link capacities which were 100 Mbps for this topology.

In our first experiment we generated 12 TCP flows, one from each source to the corresponding destination in the mirror. We take it as a basic step to show the accuracy of our mechanism which does not need any synchronization between the sending time of senders and does not require packets of different flows to be adjacent, which was required in related work.

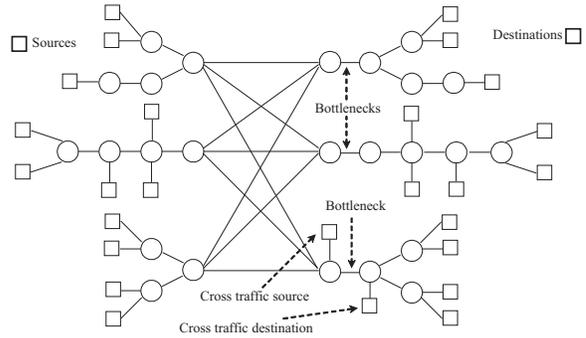


Fig. 7. Topology without common end points

We randomly start all the flows during the first second of the simulation. After gathering the delay measurements, applying SVD and then detecting the outliers to form the groups of flows which share a bottleneck, we found the Accuracy Index (*AI*) to always be 1. This is also shown in figure 8a which indicates that the first three flows, next five flows, and last four flows share a bottleneck.

As a next step, we again generated 12 TCP flows but from random sources to random destinations; again, our SVD based mechanism resulted in a very nice rank reduction and then clearly showed the outliers to form the group of flows which share a bottleneck. The results are shown in figure 8b.

As a final step, we generated three TCP flows from every sender to a randomly selected destination in each branch of the bottleneck. In this way, we generated 36 flows in total, generating foreground traffic from each sender to every branch

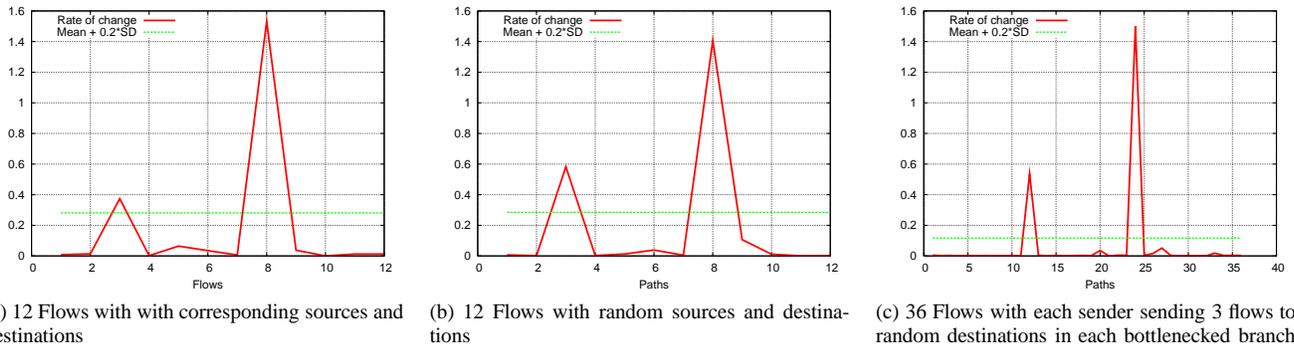


Fig. 8. Results of experiments without common end points

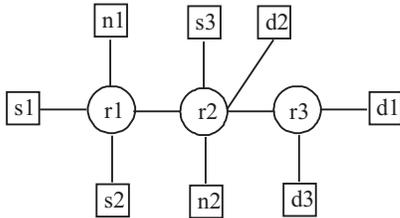


Fig. 9. Topology with a varying bottleneck

of the bottleneck. Even with such a heavy foreground and background traffic our mechanism was stable and reliable enough to produce clusters of flows having common bottlenecks with an *AI* of 1, which is shown in figure 8c.

7) *Fluctuating Bottlenecks*: It is quite possible in real networks that heavy cross traffic may cause a link to behave as a bottleneck for an end-to-end path for a while, and then a routing change or some sudden bursts may cause another link to become the bottleneck. Therefore, in our next experiments, we evaluated the adaptability of our mechanism in case of fluctuating bottlenecks.

The topology used for these experiments is depicted in figure 9. All the links have a capacity of 10 Mbps, the queue length is 250 packets, and the packet drop policy is DropTail. We generated three TCP flows as foreground flows from senders *s1*, *s2*, *s3* to the destinations *d1*, *d2*, and *d3* respectively, giving us three distinct end-to-end paths *p1*, *p2*, and *p3*. The link between *r1* and *r2* as well as the link between *r2* and *r3* may become bottlenecks. If the link between *r1* and *r2* is behaving as a bottleneck, then the paths *p1* and *p2* will share a bottleneck; if the link between *r2* and *r3* is congested, then the paths *p1* and *p3* will share a bottleneck. Hence, the bottleneck on path *p1* may fluctuate from one link to another. We congest the links between *r1*, *r2*, and *r3* by using a number of CBR flows. In addition to foreground TCP flows and CBR flows as cross traffic, we also started some Pareto flows as background traffic. These Pareto flows are originated from all sources and destinations (including nodes *n1* and *n2*) in both forward and reverse direction. The flows in forward direction were initiated with a constant bit rate of 256 Kbps whereas the flows in reverse direction were sent at a rate of 64 Kbps. The average ON and OFF times were set to 500ms and the Pareto shape parameter was set to 1.5.

In our simulation of 30 seconds we congested the first link

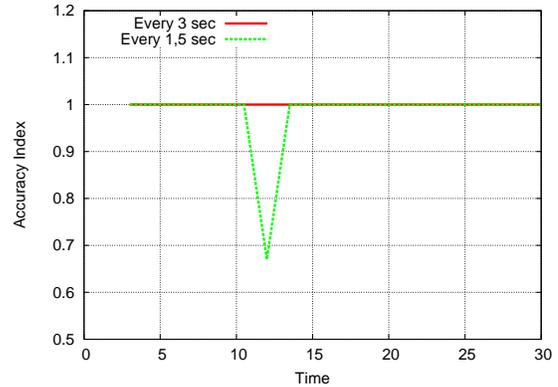


Fig. 10. Accuracy Index for every 3 seconds and for every 1.5 seconds

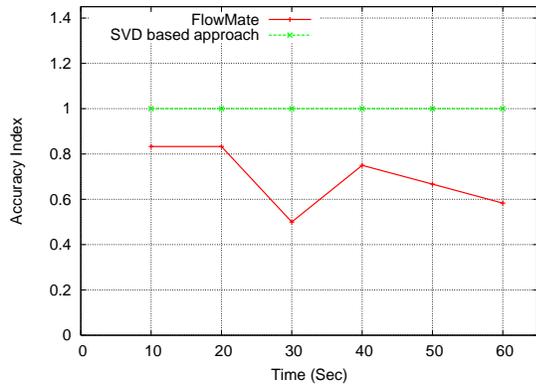
for the first 12 seconds of the simulation and then shifted the heavy cross traffic to the second link. As a result, for the first 12 seconds, paths *p1* and *p2* shared a bottleneck and for the rest of the simulation paths *p1* and *p3* shared a bottleneck. Initially we calculated the *AI* for every 3 seconds and found that it accurately detected the bottleneck shift around the 12th second of the simulation. Then, in order to check the adaptability of our mechanism at a fine grained level, we also calculated the *AI* for every 1.5 seconds from the 9th second of the simulation till the 15th second of the simulation because this was the time when the bottleneck shifted. This time we noticed a drop in the *AI* (suggesting an inaccurate clustering) around the 12th second, as shown in figure 10.

From this experiment we came to know that our mechanism was efficient enough to adopt to the shift in the bottleneck within 3 seconds, which is quite acceptable because, if the bottleneck fluctuates every time within less than 3 seconds, then this information will probably not be very useful.

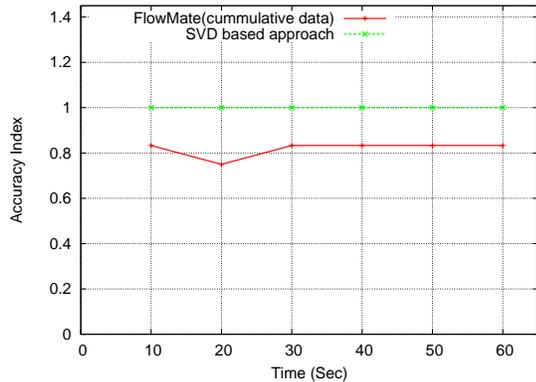
### C. Comparison with FlowMate

We compared our mechanism with FlowMate which, to the best of our knowledge, is the best of the mechanisms in terms of accuracy and computational complexity.

1) *Asymmetric Topology*: We used the asymmetric topology illustrated in figure 2b for the comparison. We used 256 Kbps Pareto flows as background traffic from the source to all destinations and 64 Kbps Pareto flows as background traffic in the reverse direction from all destinations to the source.



(a) With symmetric topology



(b) With asymmetric topology

Fig. 11. Comparison with FlowMate

All flows started randomly during the initial phase of the simulation. We used TCP flows as foreground traffic from the source to all destinations to measure the forward delays for all paths during the whole period of simulation, and these TCP flows were also started randomly during the first second of the simulation. As the FlowMate and our SVD based method both require forward delays along the paths for shared bottleneck detection, we used the same data for both methods. We applied the FlowMate algorithm and our SVD based mechanism after every 10 seconds to cluster the flows with shared bottlenecks and calculated the *AI*. Our mechanism always produced accurate results whereas the results produced by FlowMate were fluctuating under the 100% accuracy line as shown in figure 11a. Most of the time the performance of FlowMate was degraded by false splitting i.e. a correct cluster was split into further clusters and in a few occasions false sharing was also observed. This is shown in table III.

2) *FlowMate with Cumulative Data*: We used the setup from in section V-C1 to evaluate and compare FlowMate in a different manner. We decided to apply both mechanisms on accumulated data. The FlowMate is very selective in terms of picking samples to apply the correlation test as it needs adjacent packets for the cross measure  $M_x$ , and for the auto measure  $M_a$  it needs packets separated by a certain value  $t$ . So the main reason behind this experiment was to observe and compare the performance of FlowMate when it has more than the sufficient amount of data. Our SVD based mechanism was

TABLE III

RESULTS OF FLOWMATE AND OUR SVD BASED APPROACH WITH THREE CORRECT CLUSTERS ( $\{1, 2, 3\}$ ,  $\{4, 5, 6, 7, 8\}$ ,  $\{9, 10, 11, 12\}$ ). OUR SVD BASED APPROACH ALWAYS REPORTED THE ACCURATE CLUSTERS WHEREAS FALSE SHARING AND WRONG SPLITTING WAS OBSERVED IN FLOWMATE RESULTS.

Output Clusters	AI	Interpretation
$\{1, 2\}$ , $\{3\}$ , $\{4, 5, 6, 8\}$ , $\{7\}$ , $\{9, 10, 11, 12\}$	0.833	2 errors: false splitting
$\{1, 2, 3\}$ , $\{4, 5, 6\}$ , $\{7, 8\}$ , $\{9, 10, 11\}$ , $\{12\}$	0.833	2 errors: false spitting
$\{1, 2, 3, 4, 5, 6, 7, 8, 10\}$ , $\{9, 11, 12\}$	0.500	6 errors: false sharing
$\{1, 2, 3, 4\}$ , $\{5, 6, 11\}$ , $\{7, 8\}$ , $\{9, 10, 12\}$	0.750	3 errors: 2 false splitting and 1 false sharing
$\{1, 2, 12\}$ , $\{3, 11\}$ , $\{4, 5, 6, 7, 8, 10\}$ , $\{9\}$	0.667	4 errors: 1 false splitting and 3 false sharing
$\{1, 2, 10\}$ , $\{3\}$ , $\{4, 5, 6\}$ , $\{7\}$ , $\{8, 11\}$ , $\{9, 12\}$	0.583	5 errors: 3 false splitting and 2 false sharing

already proven to be good enough to provide accurate results with a small amount of data, and this experiment also showed that our SVD based mechanism can perform well even smaller amount of data. We calculated *AI* after every 10 seconds of the simulation time after applying both mechanisms on accumulated data. FlowMate performed better than before and resulted in improved results as shown in figure 11b but it still could not reach the 100% accuracy level. Another thing which observed in this experiment was the nature of errors committed by FlowMate for accumulated data the errors were only due to false splitting and we did not observe any false sharing.

## VI. CONCLUSION AND FUTURE WORK

In this paper we introduced a novel method for detecting shared bottlenecks in the network. By using SVD and applying a scheme for outlier detection on its output, we managed to obtain a perfectly accurate result in almost all of our simulations. The only exceptions from this rule are the obvious case of RED (which works against our mechanism) and the simulation of varying bottlenecks, where we deliberately tried to find the limit of our mechanism's ability to adapt to dynamic network conditions. In the RED simulations, we were able to enhance the precision by projecting the diagonal matrix over second column of the left singular matrix.

Clearly, this high reliability is the main strength of our mechanism. We believe that, combined with its operational simplicity, we may have reached a turning point with our scheme where shared bottleneck detection may no longer be a mere measurement exercise but can in fact become a key practical element of future network or transport layer mechanisms. In addition to extending the performance evaluation of our mechanism with real life measurements (here, some preliminary results obtained with the "Emulab" testbed [14] already exist, and they are similar in precision to the simulation results presented in this paper), we intend to contribute to this change in several ways in our future work, where we will embed our shared bottleneck detection method in the transfer delay prediction system that was briefly mentioned in the introduction as an immediate next step.

## VII. ACKNOWLEDGMENTS

We would like to thank Marek Wiczorek, Yehia El Khatib, and Werner Heisz for their contributions.

## REFERENCES

- [1] D. Rubenstein, J. Kurose, and D. Towsley, "Detecting shared congestion of flows via end-to-end measurement," *IEEE/ACM Transactions on Networking*, vol. 10, no. 3, pp. 381–395, June 2002.
- [2] M. S. Kim, T. Kim, Y. Shin, S. S. Lam, and E. J. Powers, "A wavelet-based approach to detect shared congestion," in *ACM SIGCOMM 2004*, August 2004.
- [3] O. Younis and S. Fahmy, "Flowmate: Scalable on-line flow clustering," *IEEE/ACM Transactions on Networking*, vol. 13, no. 2, pp. 288–301, April 2005.
- [4] D. Katabi, I. Bazzi, and X. Yang, "A passive approach for detecting shared bottlenecks," in *10th IEEE International Conference on Computer Communications and Networks*, October 2001.
- [5] M. S. Kim, T. Kim, Y. Shin, S. S. Lam, and E. J. Powers, "Scalable clustering of internet paths by shared congestion," in *IEEE INFOCOM 2006*, April 2006.
- [6] K. Harfoush, A. Bestavros, and J. Byers, "Robust identification of shared losses using end-to-end unicast probe," in *8th IEEE International Conference on Network Protocols*, November 2000.
- [7] G. H. Golub and F. V. Loan, "Matrix computations," in *Johns Hopkins University Press*, 1996.
- [8] S. Camtepe, M. Krishnamoorthy, and B. Yener, "A tool for internet chatroom surveillance," *Intelligence and Security Informatics*, 2004.
- [9] S. Chandrasekaran, B. S. Manjunath, Y. F. Wang, J. Winkler, and H. Zhang, "An eigenspace update algorithm for image analysis," *Graph. Models Image Process.*, vol. 59, no. 5, pp. 321–332, 1997.
- [10] ns Network Simulator: <http://www.isi.edu/nsnam/ns/>.
- [11] S. Floyd and V. Jacobson, "Traffic phase effects in packet-switched gateways," *SIGCOMM Comput. Commun. Rev.*, vol. 21, no. 2, pp. 26–42, 1991.
- [12] "RED parameters. <http://icir.org/floyd/red.html#parameters>."
- [13] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Trans. Netw.*, vol. 1, no. 4, pp. 397–413, 1993.
- [14] "The emulab testbed. <http://www.emulab.org>."



**Bülent Yener** Bülent Yener is an Associate Professor in the Department of Computer Science and Co-Director of Pervasive Computing and Networking Center at Rensselaer Polytechnic Institute in Troy, New York. He is also a member of Griffiss Institute for information assurance. He received MS. and Ph.D. degrees in Computer Science, both from Columbia University, in 1987 and 1994, respectively. Before joining RPI, he was a Member of the Technical Staff at the Bell Laboratories in Murray Hill, New Jersey. His current research interests include bioinformatics, medical informatics, routing problems in wireless networks, Internet measurements, quality of service in the IP networks, and the Internet security. He has been serving on the Technical Program Committees of leading IEEE conferences and workshops. Currently his is an associate editor of ACM/Kluwer Winet journal and the IEEE Network Magazine. He is a Senior Member of the IEEE Computer Society and the Chair of IEEE TCCC and TAC.



**Muhammad Murtaza Yousaf** Muhammad Murtaza Yousaf is a PhD student at University of Innsbruck, Austria. He completed his M.Sc. with distinction in Computer Science from University of the Punjab, Lahore, Pakistan. His research interests include network measurements, scheduling, Grid computing and computational theory.



**Michael Welzl** Michael Welzl passed his Ph.D. defense at the University of Darmstadt / Germany with distinction in November 2002, and received his habilitation from the same University in June 2007. He spent two years as a research assistant at the Telecooperation department, University of Linz / Austria, before joining the faculty of the newly founded Institute of Computer Science at the University of Innsbruck / Austria in November 2001, where he now heads a research team on Network Support for Grid Computing in the context of the

Distributed and Parallel Systems group.