

PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

MMS: a multihome-aware media streaming system

Ahsan Habib, John Chuang

Ahsan Habib, John Chuang, "MMS: a multihome-aware media streaming system," Proc. SPIE 6071, Multimedia Computing and Networking 2006, 607106 (16 January 2006); doi: 10.1117/12.642384

SPIE.

Event: Electronic Imaging 2006, 2006, San Jose, California, United States

MMS: A Multihome-aware Media Streaming System*

Ahsan Habib
Siemens Technology to Business Center
Berkeley, CA 94704
ahsan.habib@ttb.siemens.com

John Chuang
School of Information Management and Systems
University of California, Berkeley, CA 94720
chuang@sims.berkeley.edu

ABSTRACT

Multihoming provides highly diverse redundant paths in terms of average hop count, latency, loss ratio, and jitter. In this paper, we first explore topological path diversity and show that multihoming can significantly reduce the path overlap when a multihomed receiver conducts media streaming from a set of suppliers. We then design a multihome-aware media streaming system (MMS) that exploits topological path diversity by splitting a streaming session over the available physical links to reduce path overlap among the suppliers, and migrating a connection from one path to another if the current path is congested. A network tomography-based monitoring mechanism is developed to identify congested path segments. Through a series of experiments in the wide area Internet, we show that multihoming provides streaming at a higher rate comparing to a single service provider. On average the quality of streaming sessions is improved by 30% or more.

1. INTRODUCTION

Digital subscriber line (DSL) and cable modem technologies are the two most widely deployed broadband residential access solutions today. Both DSL and cable modem access providers offer services with data rates in the range of 100 Kbps to a few Mbps. While these are significantly faster than a dial-up modem, they are still inadequate in supporting distributed multimedia applications such as high quality video streaming and HDTV. For a multimedia streaming application, the quality of a streaming session depends on a combination of factors, ranging from the characteristics of streaming sources to the characteristics of network paths. The goodness of a network path depends on the available bandwidth, packet loss rate, packet delay, and path overlap. A user can improve bandwidth capacity and path diversity, and thus overall path goodness, by subscribing to multiple network providers. This practice is known as *multihoming*. While multihoming has historically been limited to enterprises, the declining costs of residential access services may make residential multihoming an attractive option for many households in the near future.

Multihoming provides backup links in case of link failure and the ability to reach destinations using alternative paths. The redundancy and path diversity provided by multihoming increase reliability in network communications. Moreover, during congestion, an intelligent entity at the user side can exploit the path diversity to improve the application level performance. The performance benefits of multihoming are analyzed in the context of enterprise networks in^{7,8} which show that accessing the Web can be 25% faster using two providers than with a single provider.

In our previous study,¹¹ we show that the individual ISPs may provide similar end-to-end path dynamics on average (e.g., in terms of average hop count, average latency, or average loss ratio), the performance gap between ISPs may be significant for specific destinations. Therefore, there exists a significant opportunity for applications to improve its performance via multihoming. We also show that residential multihoming provides more than last hop (“first mile”) path redundancy. For the vast majority of source-destination pairs examined, less than 30% of hops are overlapped between the paths of a multihomed host and a destination. A streaming application can effectively utilize topological path diversity to avoid congested path segments during a streaming session.

A number of studies^{6,13,18,24} show the effectiveness of Internet streaming by exploiting path dynamics and multipath routing. Furthermore, streaming quality is strongly influenced by the choice of suppliers.¹² Thus, if available, a good quality supplier should be used unless the source-destination path is congested. If part of the path is congested, multihoming can bypass the congested path and keep providing high quality streaming without replacing the supplier.

In this paper, we first show that multihoming significantly reduces the path overlap from a receiver to a set of suppliers. We then design a Multihome-aware Media steaming System (MMS) to show how a streaming application

This work is supported in part by NSF-ITR awards ANI-0085879 and ANI-0331659.

can effectively use the path diversity due to multihoming by (i) splitting a streaming session over the alternative paths and/or (ii) migrating a connection from one path to another to avoid congestion. We develop a network tomography-based connection monitoring mechanism to identify the location of a congestion that can be avoided by connection migration. Through a series of wide-area Internet experiments, we show that connection splitting provides more than 30% improvement in quality of service (QoS) of a streaming session, and connection migration can effectively bypass a congested path. Connection splitting and connection migration together provide optimal-quality streaming unless all the available paths provided by multihoming are congested.

The rest of the paper is organized as follows: Section 2 describes the test-bed and tools we use in this work. Section 3 provides an analysis of topological path diversity when a multihomed receiver communicates with one or multiple sources. Section 4 describes the proposed MMS and how it takes advantage of multihoming to improve application level quality. We evaluate multihoming streaming in a series of experiments in Section 5. Section 6 describes some existing works in the literature that can be leveraged in multihoming streaming systems. Section 7 concludes the paper.

2. EXPERIMENTAL SETUP

A typical setup of residential multihoming is shown in Figure 1, where user *A* is connected to the Internet via *ISP*₁ and *ISP*₂. When *A* communicates with *B*, whose connection is not multihomed, there are two separate end-to-end paths via *ISP*₁ and *ISP*₂. User *C* is multihomed, and user *A* has even more alternative paths to reach *C*.

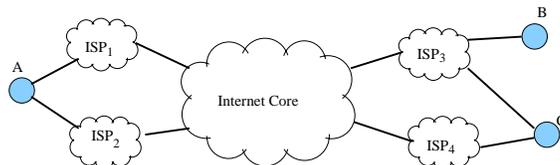


Figure 1. User *A* is multihomed by connecting to the Internet through *ISP*₁ and *ISP*₂. User *A* has two alternative paths to reach *B* and even more options to reach *C*.

2.1. Test-bed and tools

In our multihoming test-bed, a set of hosts are connected to two ISPs via an EDIMAX BR-6524² broadband router. Other choices of hardware are the Cyberpath MH200¹ multi-homing gateway, which is similar to the Edimax or the Internet Service Management Device by Rether Networks,⁴ which is a load balancing system designed for enterprise networks. Our choices of ISPs for this experiment are the largest providers of the two most prevalent broadband access technologies available today in San Francisco Bay area, namely Comcast for cable service, and SBC-Yahoo for DSL service. We refer to the DSL provider as *ISP*₁ and the cable provider as *ISP*₂ in this paper. The hosts in our test-bed have two network interface cards so that they can send packets simultaneously through both networks.

We use *traceroute* and *ping* to measure hop count, autonomous system (AS) count, latency, and packet loss, and Pathload¹⁵ to measure available bandwidth on the paths. We develop MMS based on PROMISE¹³—a peer to peer streaming system—to study the effectiveness of multihoming on video streaming applications. The video streaming experiments in Section 4 use a set of nodes in Planet-lab³ test-bed as the sources to provide content.

2.2. Performance metrics

In addition to throughput, loss, delay, and delay variation, we use the metric *quality* to quantify the performance of media streaming sessions. The quality of a streaming session is defined as:

$$Q = \sum_{i=1}^T \frac{Z_i}{T}, \quad (1)$$

where *T* is the number of packets in a streaming session and *Z_i* is a variable that takes value 1 if packet *i* arrives at the receiver before its scheduled play-out time, and 0 otherwise. The quality is different from throughput because it

considers the deadline of each packet. The parameter Q captures other performance parameters such as packet delay, packet loss, and jitter. A packet that misses its deadline is discarded and, therefore, does not contribute to the quality of a streaming session, similar to a lost packet. The system quality is defined as the average quality of all receivers in the system.

The metric Q captures packet loss as well as delayed packets and can be mapped into Mean Opinion Score (MOS), which is widely accepted as a measure of perceived video quality. Klaue et al.¹⁶ compute peak signal to noise ratio (PSNR) of frames during video transmission and map them into MOS. This mapping shows that average loss rate of 5% results into MOS score 4 (good) for 90% of the frames, and more than 50% frames experience poor quality (MOS ≤ 2) when the loss rate is 25% or higher.¹⁶ Using similar loss to quality relationship, we can show that when $Q=1$, the MOS is 4 (good) or 5 (excellent) for all frames, therefore, the perceived video quality is excellent. When $Q = 0.95$, around 10% frames experience poor or bad quality and the overall perceived video quality is very good. If $Q \leq 0.75$, almost half of the frames experience bad quality, and therefore, the overall perceived video quality is poor.

3. MULTIHOMING AND TOPOLOGICAL PATH DIVERSITY

In our previous study,¹¹ we show that the individual ISPs may provide similar end-to-end path dynamics on average in terms of average hop count, average latency, or average loss ratio, the performance gap between ISPs may be significant for specific destinations. In this paper, we explore topological path diversity when a multihomed user communicates with one or a set of remote hosts. We define two metrics Single-Source Path Overlap (SSPO) and Multiple-Source Path Overlap (MSPO) to determine the expected fraction of paths that are overlapped during single source and multiple source communications in a multihoming setup.

3.1. Single-Source Path Overlap (SSPO)

Figure 2 shows a specific instance of path diversity due to multihoming. User A at Berkeley has virtually two separate paths to reach user B at Taiwan. On one path, packets traverse through *att.net* and *verio.net* to reach *ascc.net* before reaching the destination. On the other path, packets traverse through *sbcglobal.net* and *teleglobe.net* to reach *ascc.net* before reaching B . Most of our experiments show that the $Edge_1$ and $Edge_2$ follow through different core networks $Core_1$ and $Core_2$ to reach B . We use the path diversity to design multipath multihoming streaming system. In this section, we analyze path diversity of a residential multihomed user during downloading from a set of sources.

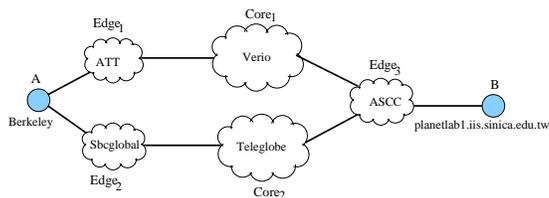


Figure 2. User A at Berkeley has two largely non-overlap paths to reach user B at Taiwan. On one path, packets traverse through *att.net* and *verio.net* to reach *ascc.net* before reaching the destination. On the other path, packets traverse through *sbcglobal.net* and *teleglobe.net* to reach *ascc.net* before reaching B .

The single source path overlap (SSPO) is the overlap between a multihomed user A and a single destination B (Figure 3a). The SSPO measures the overlap at the sender side and among the alternative paths due to multihoming. SSPO is useful in multihoming media streaming. For example, if B is a supplier of a streaming session to A and SSPO is low, B will have better chance to migrate a streaming session from one path (say via ISP_1) to another (via ISP_2) during congestion. However, if SSPO is high, it is highly likely that the shared path will be congested and B will be useless even A has two alternate paths to receive streaming. The SSPO information is also useful when user A downloads files (especially large files) from a single source. As the file transfer session lasts for a long time, it is likely that any of the paths gets congested during the session. If the congestion is not in the overlapped path segment, the user can easily migrate a session from one path to another to avoid congestion.

We compute the expected fraction of hop overlap by taking the ratio of the shared hops to the non-shared hops of the paths due to multihoming. Let H_i be the total number of hops in the path through ISP i , and E be the total

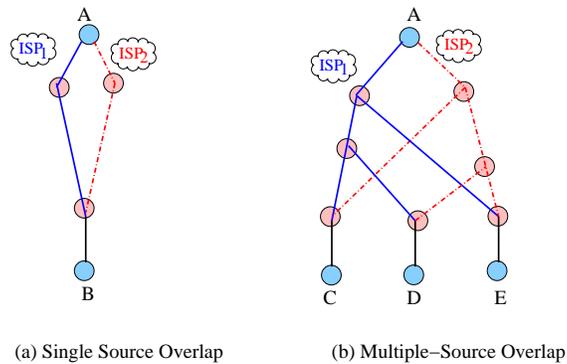


Figure 3. (a) Path overlap between a multihomed user and a destination. *A* can use any ISP to reach *B*. (b) Path overlap among a multihomed user and a set of sources. *C* and *D* have high path overlap on *ISP*₁. *D* and *E* have high path overlap on *ISP*₂. Therefore, to reduce path overlap, *C* and *E* can be assigned on *ISP*₁ and *D* on *ISP*₂ or vice versa based on actual number of hops that are overlapped in each case.

number of edges of the tree which is constructed from a destination to a multihomed source. A general definition of the expected single source path overlap *SSPO* for *k* multihoming is as follows:

$$SSPO = \frac{\sum_{i=1}^k H_i - E}{E}. \quad (2)$$

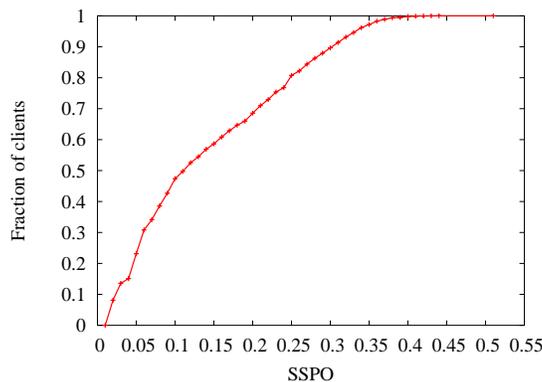


Figure 4. The CDF of SSPO for the networked hosts we studied. The expected fraction of hop overlap is less than 0.30 for 90% of the clients.

The value of SSPO varies in the range from 0 to 1, where 0 represents no overlap and 1 is 100% overlap. Figure 4 shows the expected fraction of hops that are overlapped in end-to-end paths from the network hosts to our 2-multihoming test-bed. We compute SSPO for 10,000 P2P clients that are obtained from Jian Liang, Polytechnic University at New York.¹⁷ The expected fraction of hop overlap is less than 0.30 for 90% of the clients and 0.38 for 99% of the clients. The average value of SSPO is only 0.1475. This experiment confirms that multihoming is not one hop path redundancy (90-95% overlap). Thus, two largely non-overlapped paths exist to reach a destination for a multihomed user connected to two ISPs. When one path is congested, an application can still reach a good quality source through other paths, provided that the shared path is not congested. In Section 4, we show how to detect congestion in a path segment and migrate a streaming connection from a congested path to a non-congested one.

3.2. Multiple-Source Path Overlap (MSPO)

The MSPO is useful when a user streams or downloads files (especially large files) from a set of sources. The MSPO can be used to choose the sources and distribute them on the available links such a way that the overlap among the

paths from the sources to the receiver is minimized.

Figure 3b shows the path diversity of user A to users C , D , and E . In this figure, C and D have high path overlap on ISP_1 to reach A . D and E have high path overlap on ISP_2 . Therefore, to reduce path overlap, C and E can be assigned on ISP_1 and D on ISP_2 or vice versa based on the actual number of hops that are overlapped in each case. Thus, the MSPO measures the overlap for a multihomed receiver. The MSPO can also be computed for a multihomed supplier when it acts as a supplier (instead of a receiver) of a session and wishes to know which of the available paths provides less path overlap to the receiver.

To compute the MSPO, we generate a tree from the sources to a receiver by assigning each source to a ISP. The MSPO is the expected fraction of edges of the tree that are shared by two or more paths. Let S be the set of sources that are the leaves of the tree. A subset of the sources $s_j \in S$ are assigned to ISP_j . Let H_i be the total number of hops on the path to source i , and $E_{\{s_1, s_2, \dots, s_k\}}$ be the total number of edges of the tree when a set of sources s_j is assigned on ISP_j . We compute MSPO for each possible assignment of sources to the ISPs and obtain the minimum value. The MSPO is defined for k multihoming as follows:

$$MSPO = \min_{\forall_j (j=1 \dots k), s_j \text{ on } ISP_j} \frac{\sum_{i=1}^{|S|} H_i - E_{\{s_1, s_2, \dots, s_k\}}}{E_{\{s_1, s_2, \dots, s_k\}}} \quad (3)$$

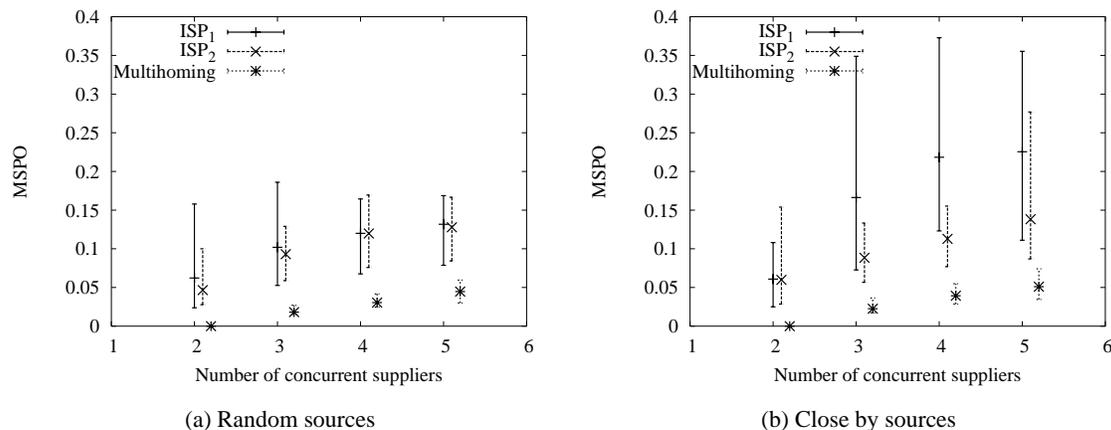


Figure 5. Multiple sources path overlap when the receiver is multihomed. Clients who are less than 100 ms RTT away from the receiver are considered as close by sources. The graph shows minimum, average, and maximum MSPO over 10 samples for each source set. The overlap is reduced at least by 2-3 times due to multihoming.

If the value of MSPO is 0, there is no overlap. A higher value represents a high overlap among the sources. The normalized value of MSPO by the number of sources is bounded by 1. We compute the MSPO metric in two different setups: (a) when the receiver is multihomed (b) when the sender is multihomed.

(a) Receiver is multihomed. Figure 5 shows the MSPO for individual ISPs and multihoming when the receiver is multihomed. The number of sources is varied on the X-axis. In Figure 5(a), the sources are chosen randomly from the 10,000 P2P clients and in Figure 5(b) the sources are chosen randomly from those who have $RTT \leq 100$ ms to the receiver. Both figures show that multihoming can reduce the overlap significantly. When there are only two sources, multihoming put them on separate ISPs, and there is no overlap ($MSPO = 0$). For a larger set of sources, on average, the overlap can be reduced by at least 2-3 times due to multihoming. The best-case results are even more promising for each supplier set. Moreover, the MSPO can be reduced further with multihoming by choosing close by sources based on RTT. Individual ISPs have higher overlap for the close by sources, and multihoming reduces the overlap significantly. Therefore, an application can use MSPO in route control decision to reduce overlap among the set of sources. Our experiments with video streaming shows that assigning sources on the ISPs based on MSPO improves application level QoS.

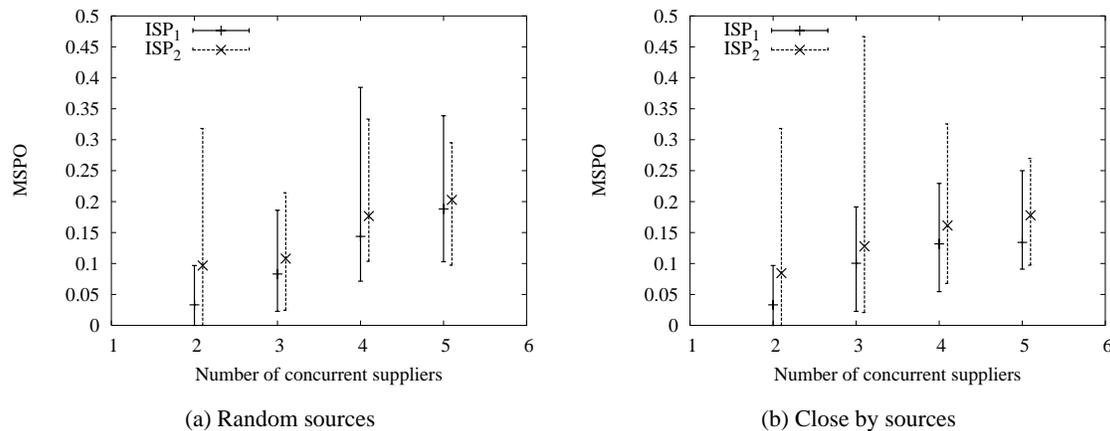


Figure 6. Multiple sources path overlap when the sender is multihomed. Using either ISP can be beneficial in some cases.

(b) Sender is multihomed. When a multihomed host wishes to be a source, it can use either of the available paths (or use both at the same time). Figure 6 shows how the multihomed source affects the expected path overlap. The average value of $MSPO$ is better for ISP_1 , however, the variation of $MSPO$ shows that taking either ISP can be beneficial in some cases.

In summary, $SSPO$ and $MSPO$ capture the path overlap on sender side and receiver side respectively. $SSPO$ determines whether a source will be able to migrate session from one path to another during congestion. $MSPO$ measures overlap among the paths from a set of sources to a destination, and multihoming can reduce the $MSPO$ by at least 2-3 times comparing to a single ISP. Therefore, a multihome-aware streaming application can use $MSPO$ to assign suppliers to the available links in order to minimize the path overlap.

4. MMS

We first describe the requirements of a streaming system, and then present a multihoming streaming model, which is used to design the proposed multihoming streaming system MMS.

4.1. Streaming system requirements

We study a streaming system to understand the requirements to provide high quality streaming. The system encompasses the key functions of object lookup, user-based aggregated streaming, and dynamic adaptation to network and supplier conditions. The quality of a supplier depends on its availability, offered rate, and outgoing link capacity. To avoid path sharing among multiple suppliers, a streaming system should leverage the underlying network topology and performance information in selecting suppliers. Techniques to cope with fluctuations in network service include forward error correction coding, multi-description coding, and sending rates adjustment. The system monitors user status to react to user and connection failure. A dynamic switching mechanism is required to replace a failed supplier without disrupting a streaming session.

4.2. Streaming model

We develop MMS on top of PROMISE,¹³ a peer-to-peer streaming system, where a client requests a media file and receives the stream in real time from a set of supplying users. The selection algorithm of PROMISE determines an *active sender set* $\mathbb{P}^{actv} = \{p_1, p_2, \dots, p_N\}$ that is likely to yield the best quality for a streaming session. We note that we can integrate any other supplier selection algorithm with our MMS.

In a MMS streaming session, N suppliers of the active set \mathbb{P}^{actv} are distributed over the available physical links. To assign the suppliers to a ISP (or link), we compute the overlap ($MSPO$) among the suppliers. The $MSPO$ determines the supplier assignment on each ISP to reduce the overall overlap. Using this ISP assignment, we obtain the topology tree for a streaming session, where receiver is the root and the suppliers are the leaves.

4.3. Path goodness

We define a *path* in terms of a supplier p and a receiver r , or a supplier and a receiver on a physical link l . We use *path goodness* to determine the quality of a path in terms of its resources and dynamics. The path goodness is used to decide whether a supplier will provide high quality streaming on a specific path. If the goodness of a path deteriorates in a streaming session, we need to reduce the streaming rate on the path or migrate the streaming session to a different path. The path goodness is computed from packet loss ratio, packet delay variation, and the number of suppliers that simultaneously share a path.

Let Z_p be a random variable that takes value 1 if packets arrive at the receiver on path $p \rightsquigarrow r$ before their scheduled play-out time, and 0 otherwise. The variable Z (also used in equation 1) takes into account both packet loss and delayed packets. We use a weight w^p to determine whether the available bandwidth $A_{p \rightsquigarrow r}$ on the path $p \rightsquigarrow r$ is greater than the rate R_p contributed by supplier p during a streaming session. If multiple suppliers share a path segment, the weight is determined considering all suppliers sharing the segment. For example, let a set of selected suppliers $S \in \mathbb{P}^{actv}$ share a path segment $i \rightarrow j$ to reach receiver r . The path segment has available bandwidth $A_{i \rightarrow j}$. Supplier $p \notin S$ shares the segment $i \rightarrow j$ with the suppliers in S to reach r . If p wants to stream at a rate R_p , the weight $w_{i \rightarrow j}^p$ for supplier p is defined as the following:

$$w_{i \rightarrow j}^p = \min \left(1, \max(0, (A_{i \rightarrow j} - \sum_{s \in S} R_s) / R_p) \right). \quad (4)$$

This is because if a segment has a bandwidth equal to or higher than the aggregate rate contributed by the suppliers S , this segment will not be congested with this aggregate rate, and its weight is set to 1. Otherwise, the weight is a fraction proportional to the remaining bandwidth along the path. We admit that measuring available bandwidth on a path is expensive, and measuring available bandwidth on individual path segment is even more expensive. The good news is that we do not need the accurate value of the available bandwidth, instead with far less overhead, we can test whether a path can accommodate the aggregated rate from the suppliers sharing the path, which perfectly serves our purpose.

Using Z and w , we define path goodness. If $G_{p \rightsquigarrow r}$ is the goodness of the path $p \rightsquigarrow r$, we express it mathematically as follows:

$$G_{p \rightsquigarrow r} = Z_{p \rightsquigarrow r} \prod_{i \rightarrow j \in p \rightsquigarrow r} w_{i \rightarrow j}^p. \quad (5)$$

We use path goodness G to compare two or more paths. A path with higher G has better path goodness. We need to measure packet loss, packet delay, and available bandwidth for each supplier on different paths to compute the path goodness. To obtain measurements from both links, we maintain data connection D_p and control connection C_p of each supplier p on separate links. We note that the measurements are done passively from the streamed data.

4.4. Multihoming streaming mechanism

Streaming in a multihoming environment assigns the suppliers on different physical links that minimizes the overall overlap and improves the expected rate at the receiver provided that the available bandwidth of each link is not exceeded. Our multihoming streaming system MMS takes advantage of multihoming environment in two different ways: (i) initially, it splits a streaming session from one physical link into multiple physical links to reduce the overlap among the suppliers to a receiver, and then, (ii) it migrates part of a streaming session from one physical link to another during congestion. Figure 7(a) shows an instance of connection split in multihoming streaming. Connections from B and C to A are split over the two physical links because the MSPO is minimum in this case. Figure 7(b) shows an instance of connection migration in multihoming streaming. Connection from C to A is migrated from the path via ISP_2 to the path via ISP_1 when the path for supplier C via ISP_2 is congested but the other path is not congested.

(i) Splitting streaming session. A streaming session streams data from multiple suppliers. Without multihoming, all connections of a streaming session are carried on the same physical link (shared path). If a shared segment $i \rightarrow j$ is congested, replacing a supplier does not improve streaming quality. The new supplier p' will share the segment $i \rightarrow j$ with other suppliers and the weight $w_{i \rightarrow j}^{p'}$ (equation 4) for the shared segment is less than one. Because the available

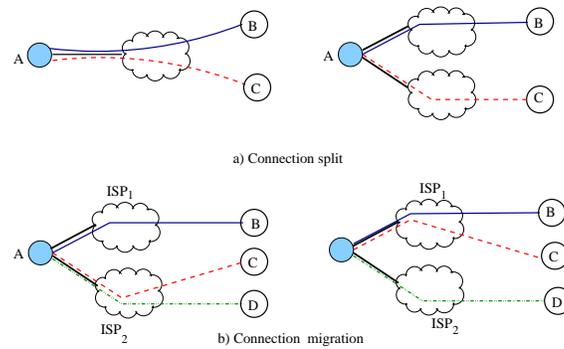


Figure 7. Connection split and connection migration with multihoming. (a) Connections from B and C to A are split over the two physical links. (b) Connection from C to A is migrated from the path via ISP_2 to the path via ISP_1 when the path for supplier C via ISP_2 is congested but the other path is not congested.

bandwidth $A_{i \rightarrow j}$ is lower than the aggregate rate of all suppliers sharing the segment. This makes the path goodness (equation 5) for the suppliers less than one and the expected rate for the receiver is reduced.

Multihoming provides non-shared physical path at the receiver end. Thus, splitting a streaming session over multiple links can reduce the likelihood of congestion. In Section 3, we show that multihoming can provide largely non-overlapped paths. Splitting connections can make the weight $w_{i \rightarrow j} = 1$ for the suppliers who were previously experiencing congestion. The connection split increases the path goodness and G_p becomes $G_{p \rightsquigarrow r} = Z_{p \rightsquigarrow r}$. However, if the congestion is at a path segment that is shared by the suppliers even after the connection split, multihoming cannot alleviate the congestion.

(ii) Migrating streaming session. Even if a streaming session uses all available paths due to multihoming, the session can experience poor quality streaming due to changes in available bandwidth on the paths. If a connection from supplier p experiences poor quality on link j , we can migrate the connection to link k , if path goodness of link k is higher than that of link j i.e. $G^j - G^k < \delta$, where δ is a small threshold to avoid unstable path switching due to a small gain. This connection migration can reduce the likelihood of congestion because the suppliers on the link j now have less competition for the available bandwidth. Thus, the connection migration can improve the goodness G of the paths from the suppliers to the receiver. If the congestion is on a path segment shared by all suppliers, the connection migration and even supplier replacement cannot improve the quality of the streaming session. In that case, we need to reduce the streaming rate.

In practice, connection migration is primarily shifting data from one path to another. Instead of closing the current connection and opening a new one, a receiver in MMS opens two concurrent data connections through both ISPs to each supplier. Based on the quality of each path, the data rate is divided over both paths. Upon detection of congestion, the data on the congested path is switched over the non-congested path.

4.5. Network tomography-based connection monitoring

We use network tomography to identify the location of congestion in end-to-end paths from a supplier to a multihomed user. Network tomography is a mechanism to infer per-link internal characteristics from end-to-end measurements. Duffield *et al.*⁹ use packet “stripes” (back-to-back probe packets) to infer link loss by computing the correlation of packet loss within a stripe at the destinations. To infer loss, a series of probe packets, called a stripe, are sent from one end host to two other end hosts with no delay between the transmissions of successive (usually three) packets. The first packet of a 3-packet stripe are sent to one host and the last two to the other host. If a packet reaches any receiver, we can infer that the packet must have reached the branch point. Using number of packets reach to the hosts, we can calculate the successful transmission probability of the links of the topology.

In MMS, we do not conduct any active probing, instead, use the streamed data to measure the congestion passively. For example, a supplier S forms a two-leaf binary tree to a multihomed receiver. We denote R_1 and R_2 are two end points at the receiver side for ISP_1 and ISP_2 respectively (Figure 8). S sends media data as a collection of stripes, i.e., out of three back-to-back packets, two go to R_1 and the third goes to R_2 or vice versa. Then, the receiver estimates

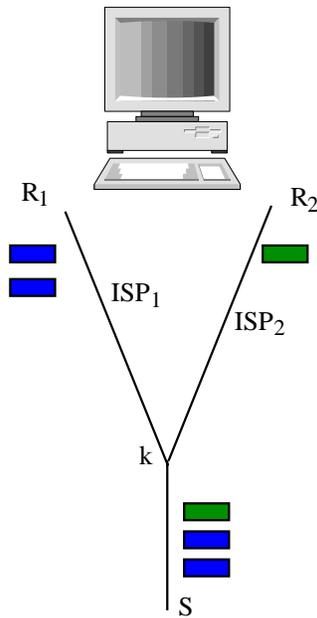


Figure 8. Inferring congestion status of links and the shared part between a streaming supplier S and a multihomed receiver R via ISP_1 and ISP_2 .

how many packets of a stripe reach via ISP_1 and ISP_2 to estimate the loss ratio on each path segment $S \rightarrow k$, $k \rightarrow R_1$ and $k \rightarrow R_2$.

MultihomingStreaming

1. Select the active set \mathbb{P}^{actv} using any supplier selection algorithm
2. Assign suppliers on available physical links that minimizes the MSPO.
3. **for** each supplier p
4. Open a data connection D_p on the assigned link and a control connection C_p on a separate physical link
5. Passively measure Z_p^l and G_p^l for each supplier p on link l from streamed data
6. **for** links j and k
7. **if** $\{ G_p^j < G_{thresh}^j \}$
8. **if** $\{ G_p^k \geq G_{thresh}^k \}$
9. Migrate data connection of supplier p on link j to link k provided $\sum_{j=1}^N R_j^k < A_k$.
10. Move its control connection to another link if both of them are on the same link
11. **else** replace supplier p on link j with a backup supplier
12. **end for**
13. **end for**

Figure 9. Streaming in a multihomed environment. The algorithm shows supplier selection, connection split, and connection migration to cope with the network dynamics and provide high quality streaming.

4.6. Multihoming streaming algorithm

Figure 9 shows the streaming algorithm in a multihomed environment. Line 1 is the supplier selection process. Any supplier selection algorithm can be plugged in our streaming system. Suppliers are split over the physical links in Line 2. For each supplier, data and control connections are open over separate physical links so that we can passively measure the packet delay and loss for the supplier over different paths (Line 4 and 5). The data connection is open on the assigned link, and the control connection that does not consume a lot of bandwidth is usually open on any other

link. The connection migration is shown in Lines 7-10 that uses the parameter G to estimate the quality of current streaming on a link and to estimate whether migrating to another link can improve quality. We use threshold G_{thresh}^l to determine whether the path goodness can provide good quality streaming on link l . The streaming application can change this value dynamically based on user response. After migrating the data connection, the control connection needs to be migrated on a different link if both of them are on the same link. If we cannot measure w at any time, the connection migration decision can be taken using Z only. If connection of supplier p experiences poor quality on all available links i.e. there exists no link for p to improve the quality, we need to replace the supplier p (Line 11).

5. MMS EVALUATION

Setup. We conduct wide area Internet experiments in the Planet-Lab³ test-bed to evaluate MMS. In our experiments, we use nodes at several U.S. universities such as Berkeley, Caltech, UCSD, Rice, Duke, MIT, Purdue, CMU, Texas, Arizona, and Stanford. We also use nodes in Germany, the UK, France, Italy, Sweden, and Taiwan. The receiver is located in the city of Berkeley and connected to the Internet via SBC (ISP_1) and Comcast (ISP_2).

We compare the performance of media streaming with or without multihoming. If a user is not multihomed, its streaming session also uses multiple suppliers to cope with network fluctuations. The performance without multihoming shown in this paper is the performance during congestion. As the streaming rate is high, the streaming application often experiences congestion when multihoming is not used. These experiments show that multihoming can effectively use the alternate paths to avoid congestion and achieve high quality streaming.

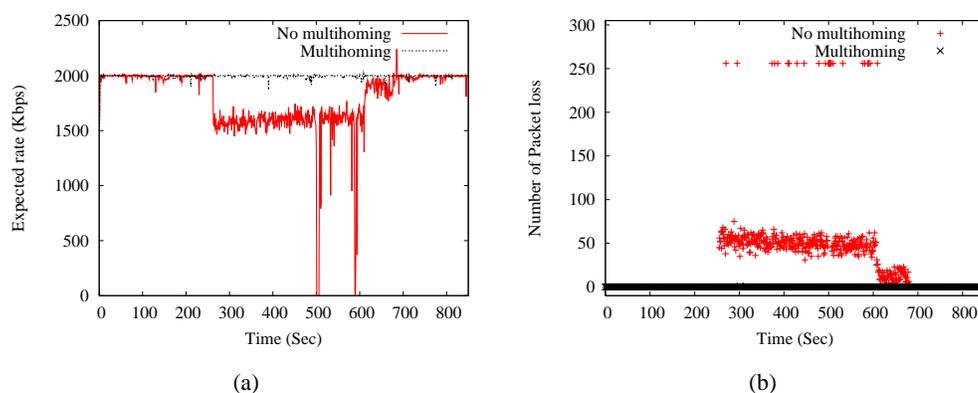


Figure 10. Expected rate in Kbps and packet loss rate per second of a streaming session with and without multihoming. The multihoming uses connection split to take advantage of available physical links. After 250 seconds, the network is congested and the rate goes down significantly without multihoming. MMS with multihoming provides high quality streaming at a steady rate.

Experiments and results. We conduct a series of experiments to show the effectiveness of connection split and connection migration provided by multihoming. First, the streaming is conducted on a single link. As ISP_2 has higher subscribed rate[†], we conduct the initial streaming on ISP_2 . The streaming rate is set at 1.5 - 2 Mbps so that ISP_2 is capable of supplying the rate. However, during congestion, the available bandwidth on the supplier to the receiver path is reduced, and the streaming session experiences poor quality. Multihoming splits the streaming over both physical links and achieves high streaming quality. Connection migration decision is taken based on the tomography-based network monitoring.

Figure 10 shows the performance of multihoming due to connection split. When there is no congestion (0 - 250 seconds), there is no significant performance gap with or without multihoming. During congestion, MMS with multihoming provides a high data rate and a low loss due to connection split (Figure 10). The data rate with multihoming has very low fluctuations, however, without multihoming the rate goes down as low as zero in a few instances. For presentation clarity, we show the rate and loss of only one specific streaming session in this figure.

[†] ISP_2 can stream at a rate as high as 3 Mbps whereas ISP_1 can stream no more than 1 Mbps.

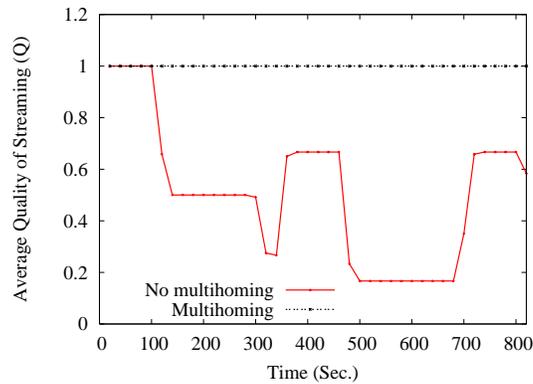


Figure 11. Average streaming quality of 6 streaming sessions with and without multihoming. The multihoming uses connection split. The stream quality degrades without multihoming because a large number of packets arrive late, and therefore missed their deadlines of play-out time. The quality improvement by multihoming is more than 30%.

Figure 11 shows the average quality of six streaming sessions. No packet missed the deadline of its play-out time when multihoming is used whereas many packets missed the deadline without multihoming. The overall quality of the streaming sessions with multihoming is excellent, and without multihoming the quality degrades substantially. This quality degradation is mainly because a large number of packets missed the play-out time deadline, and those packets are useless even though they reach the destination. Therefore, the MOS is less than 2 after 100 seconds without multihoming, whereas with multihoming the MOS is close to 5 throughout the sessions. This figure shows that the quality improvement is 30-80% with multihoming comparing to a single provider.

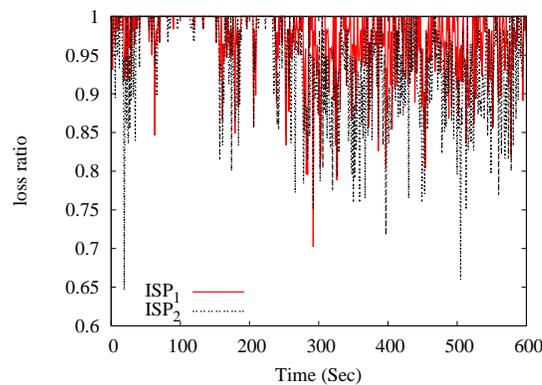


Figure 12. Inferring loss ratio using network tomography during a streaming session. The loss ratio in individual ISP paths (non-shared) varies with time. Therefore, multihoming can effectively utilize the path with low loss rates.

Connection split in a multihomed environment might not provide high quality streaming all the time. If any of the current streaming path gets congested, the congestion deteriorates the streaming quality on that path. A connection migration can avoid the congested path to achieve high quality streaming. We use both connection split and connection migration in these experiments. Initially, the streaming session is split over both links ISP_1 and ISP_2 . Three suppliers are on ISP_2 and one supplier is on ISP_1 . When one connection on ISP_2 experiences congestion (high packet loss), the connection is moved to ISP_1 . As a result, the quality of the streaming becomes significantly better. Figure 12 shows the loss inference of the non-shared paths of both ISPs for a supplier at Singapore using network tomography. This figure shows that ISP_2 experiences higher loss than ISP_1 for this supplier. Based on this loss inference, MMS migrates the connection of the supplier from Singapore from ISP_2 to ISP_1 .

Figure 13 shows that connection migration can achieve streaming with high aggregate rate and high quality. We present the rate of one specific streaming session and streaming quality of three different sessions. The average rate as

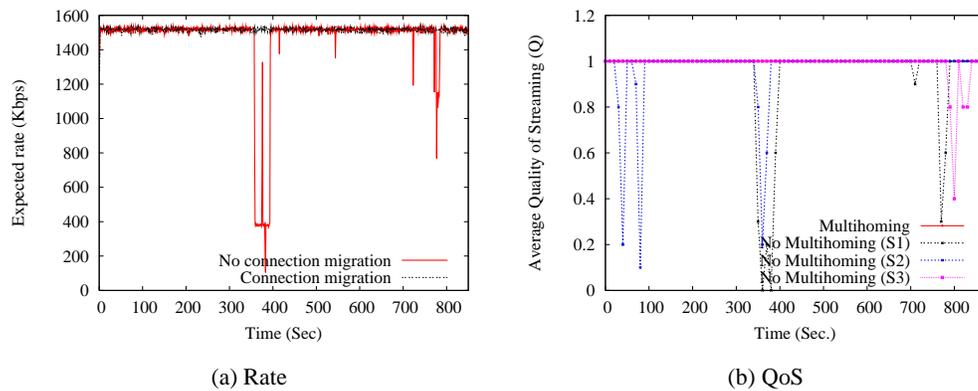


Figure 13. Average rate and quality of streaming with and without connection migration. Initially, the connections are split over both physical links. During congestion, A connection on ISP_2 is migrated to ISP_1 to achieve high quality streaming. a) The rate is shown from one specific streaming session. b) Streaming quality of three different sessions.

well as the quality is higher with connection migration compared to the setup without connection migration. Without multihoming the rate goes below the target rate several times and therefore the quality of streaming is poor.

In summary, MMS achieves significantly high quality streaming by splitting connections into multiple physical paths and by migrating connections during congestion. The connection migration decision can use network tomography-based passive network monitoring mechanism to identify the location of congestion.

6. RELATED WORK

We discuss related work on parallel downloading and multipath streaming. Each of these works has a different focus, and we can leverage some of these studies to design multihoming streaming system.

Rodrigues and Biersack²³ show that parallel download of a large file from multiple replicated servers achieves significantly shorter download time. The subsets of a file supplied by each server are dynamically adjusted based on the network conditions and the server load. This work targets bulk file transfer, not real-time media streaming.

The distributed video streaming framework^{20,21} shows the feasibility and benefits of streaming from multiple servers to a single receiver. The receiver uses a rate allocation algorithm to specify the sending rate for each server in order to minimize the total packet loss. This specification is based on estimating the end-to-end loss rate and available bandwidth between the receiver and each server. The framework does not address the selection and dynamic switching of senders.

Abdouni et al.⁶ show the benefits of multipath streaming by studying the load distribution in the context of multipath streaming. Here, a continuous media streaming session is split over multiple paths based on the path characteristics to obtain high quality streaming. This study shows that significantly higher levels of redundant information are necessary under single best path streaming in order to achieve the same level of streaming quality as multipath streaming. According to this study the frequency of losses and their correlation are important to the resulting quality of the streamed data. These observations are useful in designing multipath streaming system. As, multihoming streaming exploits the available paths to achieve high quality streaming, this study is also useful to design multihoming streaming systems.

Iyengar et al.¹⁴ develop concurrent multipath file transfer using stream control transmission protocol (SCTP) in a multihomed environment. SCTP is a transport protocol that provides a reliable transport service, ensuring that data is delivered to the receiver without error and in the same sequence as transmitted. It is necessary to investigate the effectiveness of SCTP for multihoming streaming.

Mao et al.¹⁹ design Multi-path Realtime Transport Protocol (MRTP), an application layer protocol which could use one of TCP, SCTP or UDP as transport to transmit video over ad-hoc wireless networks. MRTP specifies session

establishment and maintenance mechanisms and scheduling mechanisms over multiple paths, possibly using SCTP multihoming or UDP. This work uses multiple wireless networks to select alternate paths. Another related work for mobile networks by Phatak et al.²² where the authors propose distributing data at the network IP layer transparent to the higher layers using IP-in-IP encapsulation. Both of these applications are specially designed for wireless networks which has different loss characteristics than wired networks. The residential multihoming provides quite different network setup comparing to ad-hoc wireless networks, and we design a media streaming system that explicitly exploits the path diversity provided by residential multihoming. Our system MMS has explicit control to the route control mechanism in selecting the ISPs to improve stream quality.

The commercial products^{4,5} allow enterprises and data centers to dynamically select among upstream providers for optimal performance. Using Internet Service Management Device (ISMD) made by Rether networks, Guo et al. show that packet and flow level load balancing is equally effective for Web traffic.¹⁰ However, a simple load balancing box at the receiver end cannot ensure the quality of a streaming session. A feedback path is necessary between the streaming application and the multihoming gateway box to exchange information about the quality of the streaming. We integrate the load balancing information with the streaming application to ensure the quality of a streaming session.

Tao et al.²⁴ show that with path switching a video streaming application can improve application QoS. Our streaming system shows that path switching during congestion is beneficial for streaming application because it takes the advantage path diversity due to multihoming. Moreover, we split the connection over multiple physical links and use all links from a source to a multihomed receiver simultaneously to improve streaming quality.

7. CONCLUSION AND FUTURE WORK

Media streaming is an appealing application for residential users. Even though broadband is available in most residential households, streaming in the current Internet is still far from reality. In this paper, we have shown that a residential user can enjoy high quality streaming by utilizing multihoming that provides highly diverse paths. Multihoming significantly reduces path overlap among the suppliers of a streaming session. Our experiments in the wide area Internet shows that multihoming can improve streaming quality by 30% or higher by splitting the connections over available physical links and/or migrating connections from one path to another to avoid a congested path.

We plan to explore zero cost multihoming streaming where two users subscribe different ISPs, and then share the Internet service with each other via wireless setup. As the number of users of this wireless network is small, we believe that it will not significantly affect the quality of media streaming. Our future work also includes evaluating MMS with multihomed suppliers, which improves the upstream bandwidth constraint for residential users. Finally, it is important to understand how residential multihoming might challenge the existing architectures, business models, and industrial organization of residential access networks.

REFERENCES

1. Cyberpath inc., <http://www.cyberpathinc.com/>, 2005.
2. Edimax, <http://www.edimax.com>, 2005.
3. Planet-lab test-bed. <http://www.planet-lab.org/>, 2005.
4. Rether networks inc., <http://www.rether.com>, 2005.
5. Routsience, <http://www.routsience.com>, 2005.
6. B. Abdouni, W.C. Cheng, A. L.H. Chow, L. Golubchik, W.-J. Lee, and J. C.S. Lui. Multi-path streaming: Optimization and evaluation. In *Proceedings Multimedia Computing and Networking (MMCN '05)*, San Jose, USA, Jan. 2005.
7. A. Akella, B. Maggs, S. Seshan, A. Shaikh, and R. Sitaraman. A measurement-based analysis of multihoming. In *Proceedings ACM SIGCOMM*, Karlsruhe, Germany, Aug. 2003.
8. A. Akella, S. Seshan, and A. Shaikh. Multihoming performance benefits: An experimental evaluation of practical enterprise strategies. In *Proceedings USENIX Annual Technical Conference*, Boston, MA, June 2004.
9. N. G. Duffield, F. Lo Presti, V. Paxson, and D. Towsley. Inferring link loss using striped unicast probes. In *in Proc. IEEE INFOCOM*, April Alaska, April 2001.
10. F. Guo, J. Chen, W. Li, and T.-C. Chiueh. Experiences in building a multihoming load balancing system. In *Proceedings IEEE INFOCOM*, Hong Kong, Mar., 2004.

11. A. Habib, N. Christin, and J. Chuang. Can multihoming realistically improve end-to-end service?, May 2005. <http://www.sims.berkeley.edu/~habib/papers/path-switching.pdf>.
12. A. Habib and J. Chuang. Incentive mechanism for Peer-to-Peer media streaming. In *Proceedings International Workshop on Quality of Service (IWQoS '04)*, pages 171–180, Montreal, Canada, June, 2004.
13. M. Hefeeda, A. Habib, B. Botev, D. Xu, and B. Bhargava. PROMISE: Peer-to-peer media streaming using CollectCast. In *Proceedings ACM Multimedia '03*, Berkeley, California, Nov. 2003.
14. J. Iyengar, K. Shah, P. Amer, and R. Stewart. Concurrent multipath transfer using sctp multihoming. In *International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS'04)*, San Jose, USA, July 2004.
15. M. Jain and C. Dovrolis. End-to-end available bandwidth: measurement methodology, dynamics, and relation with TCP throughput. *IEEE/ACM Transaction on Networking*, 11(4):537–549, 2003.
16. J. Klaue, B. Rathke, and A. Wolisz. Evalvid - a framework for video transmission and quality evaluation. In *Proc. of 13th International Conference on Modelling Techniques and Tools for Computer Performance Evaluation*, Urbana, Illinois, USA, September 2003.
17. J. Liang, R. Kumar, and K. W. Ross. The KaZaA overlay: A measurement study, Sept. 2004.
18. D. Loguinov. Adaptive scalable internet streaming. Ph.D. Dissertation, the City University of New York, 2002.
19. S. Mao, S. Lin, S. S. Panwar, Y. Wang, and E. Celebi. Video transport over ad hoc networks: Multistream coding with multipath transport. *IEEE journal on selected areas in communications*, 21(10):1721–1737, Dec. 2003.
20. A. Nguyen and A. Zakhor. Distributed video streaming with forward error correction. In *Proceedings Int'l Packetvideo Workshop (PV '02)*, Pittsburgh, USA, April 2002.
21. A. Nguyen and A. Zakhor. Distributed video streaming over Internet. In *Proceedings Multimedia Computing and Networking (MMCN '02)*, San Jose, USA, Jan. 2002.
22. D. S. Pathak, T. Goff, and J. Plusquellic. Ip-in-ip tunneling to enable the simultaneous use of multiple ip interfaces for network level connection striping. *Computer Networks*, 43(6):787–804, Dec. 2003.
23. P. Rodriguez and E. Biersack. Dynamic parallel access to replicated content in the internet. *IEEE/ACM Transactions on Networking*, 10(4):455–465, Aug. 2002.
24. S. Tao and R. Guerin. Application-specific path switching: A case study for streaming video. In *Proceedings of ACM Multimedia*, New York, USA, October 2004.